IBM Content Manager OnDemand for iSeries
Common Server



Administration Guide

Version 5 Release 2

SC27-1161-01

IBM Content Manager OnDemand for iSeries
Common Server



Administration Guide

Version 5 Release 2

SC27-1161-01

Note
Before using this information and the product it supports, read the information in Appendix H, "Notices" on page 187.

Second Edition (September 2002)

This edition applies to IBM® Content Manager OnDemand for iSeries™ Common Server, Version 5 Release 2 and to all subsequent releases and modifications until otherwise indicated in new editions. This edition replaces SC27-1161-00.

© Copyright International Business Machines Corporation 2001, 2002. All rights reserved.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About IBM Content Manager OnDemand for iSeries Common Server Administration Guide (SC27–1161) . . . vii
Who should read this book vii
How this book is organized vii
Conventions and terminology used in this book viii
Prerequisite and related information ix
| iSeries Navigator. ix
How to send your comments ix
Product support x

| **Summary of changes . . . . . xi**

Part 1. Database administration . . . 1

Chapter 1. Database concepts 3
System 3
Instance 3
Database. 3
Table 4
Index 4
Journals and journal receivers 4

Chapter 2. Database maintenance . . . 5
Before you begin 5
Expiring index data 5
How to expire index data 6
Migrating indexes 7
How to migrate indexes 7

Chapter 3. Migrating and importing index data 9
Configuring the system 9
System Log messages 9
System Log user exit program 9
Archive Storage Manager 10
Storage sets 10
Application groups 10
What happens when a user queries migrated data 10
Message to the user. 11
Message to the System Log 11
Importing index data 11
Verify database storage space 11
Verify database log file space 11
Run the STRIMPOND command 11
After you import index data 11
Expiring imported migrated indexes 11
Configuring index migration 12

Part 2. Storage administration . . . 13

Chapter 4. Document storage 15
Defining document storage management 15
Application groups 15
Disk storage manager 16
Archive storage manager 16
Migrating documents 17
Migration processing in the system log 18
Removing documents 18
Removing documents from disk storage 19
Removing documents from archive storage . . . 20

Chapter 5. Backup and recovery. . . . 21
Backup considerations 21
Recovery considerations 22
Reports 22

Part 3. Using iSeries Navigator for OnDemand administration 23

Chapter 6. Installation 25
Hardware 25
Software 25
Memory 25
Disk space. 25
| Installing the iSeries Navigator interface to
| OnDemand 26
Installing the administrative client. 26
Running Setup 26
To install on a user's PC 26
To install on a network file server 27
To use automated install 27
To uninstall 27

Chapter 7. Introduction. 29
Getting started 29
Using online help 30
Adding a server 30
Logging on a server 31
Changing passwords 31
System parameters 32
Maximum Password Age 32
Minimum Password Length 32
Inactivity Time Out. 32
System Logging 33
User Exit Logging 33
Login Processing 33
Setting system parameters 34
Adding items to a server 35
New command 35
Copy command 35
Export command 36
Drag and drop operation 36
Report Wizard 36
Starting the Report Wizard 37
Using the Report Wizard 37

© Copyright IBM Corp. 2001, 2002

Chapter 8. Concepts 41
Migration Policies 41
Monitor Definitions. 41
Tape devices 42
Tape volumes. 42
Optical storage groups. 43
Disk pool storage groups 43
Users 43
User types. 44
Authority 45
Groups 46
Printers. 46
Storage sets 46
Application groups 47
Applications 47
Folders 48
About application groups, applications, and folders 48
OnDemand permissions 49
Folder permissions 50
Application group permissions 50
Specifying permissions 51
Hints and tips 54

Chapter 9. Examples 55
System configuration 55
Adding tape devices 56
Adding tape volumes 56
Creating optical storage groups. 57
Adding optical volumes 58
Creating disk pool storage groups 59
Creating monitor definitions. 60
Creating migration policies 63
Adding users. 65
The basics 65
Examples 66
Adding groups 68
Choose a server 68
Two ways to add a group 69
Adding users. 69
Adding the group 69
Examples 69
Adding server printers 71
Choose a server 71
Two ways to add a server printer 71
Adding the server printer 72
Examples 72
Adding storage sets 72
Adding a report 73
An example 73

Part 4. Loading data 95

Chapter 10. Loading spooled file data 97
Overview 97
Preparing to load reports 98
Storage space. 98
Defining the application group 98
Defining the application 98
Loading reports 99
Processing the input data 100
Loading index data 101
Loading storage objects 101
Verifying processing 102
Backing up databases. 103

Chapter 11. Loading image files . . . 105
Overview. 105
Defining the application group 105
Database Organization 105
Expiration Type 106
Permissions 106
Field Definition. 106
Defining the application. 106
Application Group 106
Data Format. 107
Indexer 107
Data Compression. 107
Defining the folder 107
Application Group 107
Permissions 107
Field Definition. 107
Field Mapping 108
Acquiring the images. 108
Accessing the image files 108
Creating index data 108
Configuring the ARSLOAD program 109
Processing the input data 110
Processing index data 110
Processing the image files 110
Verifying processing 111

Chapter 12. Loading user-defined data 113
Defining the application group 113
Database organization 113
Expiration type. 114
Permissions 114
Field definition 114
Defining the application. 114
Application Group. 115
Data Format. 115
File Extension 115
Indexer 115
Data Compression. 115
Defining the folder 115
Application Group. 115
Permissions 116
Folder fields. 116
Field Mapping 116
Accessing the input files. 116
Creating the index data 116
Configuring the ARSLOAD program 117
Processing the input data 118
Processing index data 118
Processing the Lotus WordPro files 118
Verifying processing 119

Chapter 13. Restarting a load process 121

Chapter 14. Deleting a report. 123

Part 5. Server administration . . . 125

Chapter 15. Managing the server . . . 127
Errors and alerts . . . 127
System logging facility . . . 127
Searching for and viewing messages. . . . 127
System Log user exit . . . 128
Monitoring users . . . 128
Reviewing a CD-ROM creation manifest . . . 129
How to find the server job and its attributes . . . 129
Restarting journaling . . . 130

Part 6. Appendixes . . . 131

Appendix A. Command reference . . . 133
OnDemand server commands . . . 133
ADDRPTOND . . . 133
ENDMONOND . . . 133
FNDKEYOND . . . 133
PRTRPTOND . . . 133
PRTTXTOND . . . 134
RMVRPTOND . . . 134
STRASMOND . . . 134
STRDSMOND . . . 134
STRIMPOND . . . 135
STRMONOND . . . 135

Appendix B. 5250 host connection to client viewer . . . 137
Operational and environmental considerations . . . 137
Workstation installation tasks . . . 137
Server Configuration . . . 139

Appendix C. Server printing and faxing. . . . 141
Using an OS/400 printer file to define server print parameters . . . 141
Server fax setup . . . 142

Appendix D. API and user exit reference . . . 145
API reference . . . 145
Using quotes when executing the APIs . . . 145
Using the QSHELL environment . . . 145
Calling QSHELL commands from an OS/400 command line . . . 146
How to read a syntax diagram . . . 147
ARSDATE . . . 149
ARSDOC . . . 153
ARSLOAD . . . 171
User exit reference. . . . 177
Output queue monitor user exit program . . . 177
Facsimile user exit program . . . 178

Appendix E. Automating ARSLOAD data loading . . . 179
ARSLOAD . . . 179
Automating ARSLOAD . . . 179

Appendix F. Alternative to starting the administrative client . . . 181
Starting the Administrator . . . 181
Administrator start up parameters . . . 181

Appendix G. Accessibility features 185
Keyboard input and navigation . . . 185
Keyboard input. . . . 185
Keyboard focus. . . . 185
Features for accessible display . . . 185
High-contrast mode . . . 185
Font settings. . . . 185
Non-dependence on color . . . 185
Alternative alert cues . . . 186
Compatibility with assistive technologies . . . 186
Accessible documentation . . . 186

Appendix H. Notices . . . 187
Trademarks . . . 189

Index . . . 191

About IBM Content Manager OnDemand for iSeries Common Server Administration Guide (SC27–1161)

The information in this book can help administrators maintain an IBM Content Manager OnDemand for iSeries Common Server Version 5 Release 2 (OnDemand) system.

This book describes how to:

- Maintain the database
- Manage disk and archive storage
- Use the administrative client to define reports to the system and maintain other types of objects on the system
- Load reports and other types of data into the system
- Use administrative commands to complete other types of tasks

Before you begin, we recommend that you familiarize yourself with concepts and terminology used in this book. See the *IBM Content Manager OnDemand for iSeries Common Server: Planning and Installation Guide* for information about OnDemand.

Who should read this book

This book is of primary interest to administrators who are responsible for implementing and maintaining an OnDemand system. Some administrators can use this book and the tools described in it to define reports to the system. Other administrators can use this book and the tools described in it to maintain users, groups, printers, and so forth. Still other administrators can use the administrative commands described in this book to maintain the database, disk and archive storage, extract documents from the system, and so forth.

How this book is organized

Part 1, "Database administration" on page 1 describes the key database objects that you will need to understand to perform the tasks in this book, describes how to maintain the database, and provides information about backup and recovery. This part contains the following sections:

Chapter 1, "Database concepts" on page 3

Chapter 2, "Database maintenance" on page 5

Chapter 3, "Migrating and importing index data" on page 9

Part 2, "Storage administration" on page 13 contains information about managing disk and archive storage devices and objects. This part has the following section:

Chapter 4, "Document storage" on page 15

Chapter 5, "Backup and recovery" on page 21

Part 3, "Using iSeries Navigator for OnDemand administration" on page 23 contains information about iSeries Navigator, a program that you can use to define and maintain objects on the system. For example, you can use iSeries Navigator to define reports to OnDemand and maintain OnDemand users, groups and printers. This part contains the following sections:

© Copyright IBM Corp. 2001, 2002

Chapter 6, “Installation” on page 25 describes the hardware, software, memory, and disk space requirements for the administrative client and how to install the administrative client

Chapter 7, “Introduction” on page 29 provides an introduction to the iSeries Navigator, including how to start the program, log on to a server, and maintain passwords, describes the types of tasks that you can do with the iSeries Navigator, and contains information about the system parameters that you can maintain with the iSeries Navigator

Chapter 8, “Concepts” on page 41 provides an overview of the OnDemand objects: users, groups, printers, migration policies and storage sets, application groups, applications, and folders.

Chapter 9, “Examples” on page 55 provides step-by-step instructions for using the iSeries Navigator to define users, groups, printers, migration policies and storage sets, and reports to the system

Part 4, “Loading data” on page 95 provides information about and examples of loading reports on the system. This part contains the following sections:

Chapter 10, “Loading spooled file data” on page 97

Chapter 12, “Loading user-defined data” on page 113

Chapter 11, “Loading image files” on page 105

Chapter 13, “Restarting a load process” on page 121

Chapter 14, “Deleting a report” on page 123

Part 5, “Server administration” on page 125 describes errors and alerts, system logging and how to find the server job system logging, how to find the server job, and how to restart journaling.

Appendix A, “Command reference” on page 133 contains reference information about the OnDemand server commands.

Appendix B, “5250 host connection to client viewer” on page 137 describes the OnDemand 5250 Host Connection, which allows an OS/400® application to send information to the OnDemand client workstation viewer program (the viewer).

Appendix C, “Server printing and faxing” on page 141 describes how to define some of the server print parameters in your Common Server application definition.

Appendix D, “API and user exit reference” on page 145 describes the OnDemand APIs.

Appendix E, “Automating ARSLOAD data loading” on page 179 describes how to automate the ARSLOAD command.

Appendix F, “Alternative to starting the administrative client” on page 181 describes a different way to start the administrative client.

Conventions and terminology used in this book

The term Windows® client refers to the OnDemand client software that runs under Windows 2000, Windows 98, Windows NT® 4.0 with SP5 or later, and Windows XP.

Prerequisite and related information

Use the IBM iSeries Information Center as your starting point for looking up iSeries technical information. You can access the Information Center two ways:

- From the following Web site:

  `http://www.ibm.com/eserver/iseries/infocenter`
- From CD-ROMs that ship with your Operating System/400® order:

  iSeries Information Center, SK3T-4091-02. This package also includes the PDF versions of iSeries manuals, *iSeries Information Center: Supplemental Manuals*, SK3T-4092-01, which replaces the Softcopy Library CD-ROM.

The IBM iSeries Information Center contains advisors and important topics such as CL commands, system application programming interfaces (APIs), logical partitions, clustering, Java™, TCP/IP, Web serving, and secured networks. It also includes links to related IBM Redbooks™ and Internet links to other IBM Web sites such as the Technical Studio and the IBM home page.

iSeries Navigator

IBM iSeries Navigator is a powerful graphical interface for managing your iSeries servers. iSeries Navigator functionality includes system navigation, configuration, planning capabilities, and online help to guide you through your tasks. iSeries Navigator makes operation and administration of the server easier and more productive and is the only user interface to the new, advanced features of the OS/400 operating system. It also includes Management Central for managing multiple servers from a central server.

For more information on iSeries Navigator, see the IBM iSeries Information Center.

How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information. If you have any comments about this book or any other iSeries documentation, fill out the readers' comment form at the back of this book.

- If you prefer to send comments by mail, use the readers' comment form with the address that is printed on the back. If you are mailing a readers' comment form from a country other than the United States, you can give the form to the local IBM branch office or IBM representative for postage-paid mailing.
- If you prefer to send comments by FAX, use either of the following numbers:
  - United States, Canada, and Puerto Rico: 1-800-937-3430
  - Other countries: 1-507-253-5192
- If you prefer to send comments electronically, use one of these e-mail addresses:
  - Comments on books:

    RCHCLERK@us.ibm.com
  - Comments on the iSeries Information Center:

    RCHINFOC@us.ibm.com

Be sure to include the following:

- The name of the book or iSeries Information Center topic.
- The publication number of a book.
- The page number or topic of a book to which your comment applies.

Product support

Product support is available on the Web. Click **Support** from the product Web site at:

`http://www.ibm.com/software/data/ondemand/`

The IBM support center maintains product updates for OnDemand. You can obtain the latest product update for the end-user client or parts of the administrative client software from IBM service on the Web at:

`ftp://service.software.ibm.com/software/ondemand/fixes/v71`

The current list of OnDemand server and iSeries Navigator component PTFs is available in informational APAR II13168.

If you encounter problems or errors running any of the OnDemand programs, you can call the IBM support center to obtain software problem and defect support.

Summary of changes

This edition of *IBM Content Manager OnDemand for iSeries Common Server: Administration Guide* contains new technical information. There may be some instances where changes were made, but change bars are missing. Significant changes to note are:

You can automate the loading of non-spooled file data such as PC files in IFS with the Start Monitor for OnDemand (STRMONOND) command using *DIR (directory) for the TYPE parameter. See Appendix A, "Command reference" on page 133 for information on the STRMONOND command.

Additional keywords have been added to many OnDemand commands to more precisely identify the spooled file that the command will use. The new keywords correspond to the same new keywords available for OS/400 spooled file commands, allowing you to specify the system on which the spooled file was created, as well as the spooled file creation date and time.

Portable Application Solutions Environment (PASE), a product option of OS/400, is now an optional software prerequisite for the OnDemand Common Server. PASE is required if you plan to use the new OnDemand Common Server text search function for AFPDS documents. It is also possible that, in the future, other new functions of OnDemand may require PASE.

Additions and enhancements have been made to the sample programs for both Common Server and Spool File Archive. Sample programs for Common Server can be found in QSAMPLES2 source file in library QRDARS. Sample programs for Spool File Archive can be found in QSAMPLES source file in library QRDARS.

Record Archive provides commands and application programming interfaces (APIs) that let you store and retrieve data records on optical media for users who only require occasional access to historical data. At Version 5 Release 2, this product option is provided for existing Record Archive customers to use, but there are no planned enhancements. Documentation can be found in OnDemand publications from previous releases. Please talk to your software provider about other options, such as compressed DASD.

© Copyright IBM Corp. 2001, 2002

Part 1. Database administration

This part is of primary interest to administrators responsible for configuring and maintaining the OnDemand database. This part is organized as follows:

- Database concepts
- Database maintenance
- Migrating and importing index data

© Copyright IBM Corp. 2001, 2002

Chapter 1. Database concepts

This section provides an overview of the key database objects:

- System
- Instance
- Database
- Table
- Index
- Journals and Journal Receivers

System

A database *system* represents a physical machine that contains a copy of the database manager. In OnDemand, the *server* is the physical machine referred to as the *system*.

Instance

An OnDemand instance is a logical server environment consisting of a server and its own separate database and disk space. Each OnDemand instance (the server, database and disk):

- Has its own definition of folders, application groups, applications and printers
- Must run in a single CCSID
- Has different security (users, groups, folder and application group permissions)
- Must have its name specified on commands if it is not the default instance
- Has its own System Log

Some reasons you might have multiple instances on a machine are:

- To have distinct test and production environments
- To have databases using different CCSIDs

For OnDemand, we recommend that you name your primary production, or only, instance *QUSROND*. This is the default instance that is created when you install OnDemand.

Database

A collection of data that is stored in tables. In OnDemand, generally speaking, there are two types of tables:

- System tables, which contain information about the objects you define to the system, such as users, groups, application groups, applications, folders, storage sets, and printers. There are also system tables that contain information OnDemand uses to control and maintain the system.
- Application group tables, which contain the index data for the reports that you load on the system

.

© Copyright IBM Corp. 2001, 2002

Table

A table consists of data logically arranged in columns and rows. For example, when you create an application group, the system creates a table definition that contains one column for each field that you define. When you load a report into an application group, the system adds one row to an application group table for each document contained in the report.

Index

In OnDemand, an index is a key that points to a document. An index allows more efficient access to documents by creating a direct path to a document through pointers.

You define indexes when you create an application group. The indexes should contain information that uniquely identify a document, such as date, account number, and customer name. Indexes are populated by values extracted from a report when you load a report on the system. Each row in an application group table identifies one document.

However, keep in mind that you do not want lots of indexes on a report just to have indexes. You should have a good business reason to have an index. While indexes can help you find documents faster, having too many of them can slow you down when you load reports on the system. Every time you add a new row (document) to a table, you have to add a row to each and every one of the indexes for that table. So the more indexes that you have, the longer it may take when you load a report.

The SQL optimizer automatically chooses the most efficient way to access data in tables. The optimizer takes indexes into consideration when determining the fastest access path to data.

Journals and journal receivers

Each database includes recovery journals and journal receivers, which are used to recover from application or system errors. In combination with database backups, journals are used to recover the consistency of the database right up to a point in time when an error occurred.

All instances have journals associated with them. These journals keep records of database changes. Journals prevent a failure (system power, application error) from leaving a database in an inconsistent state. They restore the state of a database to the point before the change. Journals can also enable forward recovery to any point in time before the failure.

In planning for disaster recovery, be sure to remember that journals must be stored off site, or at least safely away from the disaster, in order to recover your database beyond the point of the last full, off line backup.

Chapter 2. Database maintenance

You need to maintain the OnDemand database to keep it performing in an optimal manner. We recommend that you run the following database maintenance tasks on a regular basis:

- Expire index data that has reached its life of data and indexes period
- Migrate indexes to archive storage, if your organization needs to keep indexes for some period of time after users no longer need to retrieve the documents to which they point. This is not recommended.

Before you begin

We recommend that you do the following:

- Keep the index data for any given version of a report in the database on disk at least until such time that 99 percent of the requests for the report have passed. As a rule of thumb, if there is any chance that someone in your organization will need to retrieve a version of a report, keep the index data in the database – don't allow it to be migrated.
- Expire data periodically, and migrate index data only when absolutely necessary.
- Select Multiple Loads per Database Table when you define the Database Organization within your application group definitions. With this selection, each time that you load a report into an application group, OnDemand inserts the index records into an existing database table. Index records for every report loaded into the application group are stored in the same database table. OnDemand maintains the application group data so that, as far as a user querying the application group knows, they appear to reside in one database table. OnDemand automatically segments the application group data when it grows beyond a certain size. OnDemand maintains a segment table for each application group. The segment table provides faster query performance by limiting searches to a specific table of application group data, using a date value to construct the query. This method is the default, and should be used in most cases.

Expiring index data

Indexes *expire* (are eligible for removal) because their life of data period has passed. The indexes, and the documents that they point to, can then be removed from the system. When you remove an index, information about the document to which it points is removed from the database (the document can no longer be retrieved). However, because indexes are eligible to be removed does not mean that they will be deleted from the database. OnDemand does not delete expired index data from the database until expiration processing runs.

The application group expiration policy determines when index data is eligible for deletion from the database. You define the expiration policy when you create the application group. The following properties on the Storage Management page comprise the expiration policy:

- Life of Data and Indexes. The length of time in days to maintain index data and documents on the system. After the index data has been on the system for this number of days, it is eligible to be deleted.

© Copyright IBM Corp. 2001, 2002

Note: If you specify Never Expire, then expiration processing is disabled for the application group. (That is, index data will not be removed from the database.)

- Expiration Type. Determines whether individual indexes or an entire table of index data is deleted at a time. When OnDemand deletes index data, it either deletes a row (if the Expiration Type is Document) or deletes a table (if the Expiration Type is Segment or Load). The amount of index data in a table and the number of reports the data represents is determined by the Database Organization. If the Database Organization is Multiple Loads per Database Table, then by default, a table of index data can hold up to 10 million indexes. These types of tables usually hold the indexes for many reports. If the Database Organization is Single Load per Database Table, then each table holds the indexes for one and only one report.

A table of index data is not eligible to be deleted until the latest date in any of its rows reaches the Life of Data and Indexes period. For example, suppose that the Life Of Data and Indexes is set to 365 days, the Expiration Type is set to Segment, and the Database Organization is set to Multiple Loads per Database Table. By default, a table will contain approximately 10 million rows. Further, suppose that a report is loaded into the application group once every month and that each report adds one million rows to the database. Each table can hold the index data from approximately 40 reports. Using these assumptions, the data that is loaded into the application group in January will not be eligible to be deleted by expiration processing until 40 months (more than three years). If you need to remove the index data for a report as soon as it reaches its Life of Data and Indexes period, then set the Database Organization to Single Load per Database Table and set the Expiration Type to Segment or Load. (And run expiration processing at least once a month.)

OnDemand's disk storage manager and the archive storage manager delete the documents that expired index data points to independently of each other. OnDemand uses the application group's expiration policy to determine when indexes and documents expire and should be removed from the system. The archive storage manager marks documents for removal based on the criteria specified in the migration policy. However, you should specify the same criteria to the disk storage manager and the archive storage manager. In other words, the sum of the Duration values for all migration policy storage levels plus the value for the Days in Cache should equal the Life of Data and Indexes value in the application group.

How to expire index data

OnDemand does not delete expired index data from the database until expiration processing runs. The STRDSMOND command is the expiration utility. You can schedule the STRDSMOND command to run automatically or you can run it manually. You should make sure that the STRDSMOND command runs periodically so that OnDemand deletes indexes when it is time to do so (so that expired documents can no longer be retrieved).

When the STRDSMOND command removes indexes, it saves the following message in the System Log:

```
128 ApplGrp Segment Expire (ApplGrp) (Segment)
```

One message is saved in the System Log for each table that was deleted during expiration processing.

Note: While not recommended, if you have migrated indexes to archive media, then the STRASMOND command will perform expiration processing on that index data.

Migrating indexes

Note: This section provides an overview of the process of migrating index data from the database to archive storage. See Chapter 3, "Migrating and importing index data" on page 9 for information about configuring the system for migration processing.

OnDemand provides automatic migration to move indexes from the database to archive storage to maintain seldom used indexes for long periods of time.

Important: If you use migration to move indexes to archive storage, make sure that you migrate them after there is no longer a need to retrieve the documents to which they point.

The STRDSMOND command uses an application group's migration policy to control when migration of indexes for an application group occurs:

- Migration of Indexes. If you specify No Migration, then migration of indexes is disabled for the application group. (That is, index data will not be migrated.) If you specify Migrate After *n* Days, then index data is eligible to be migrated after reaching the specified number of days. Indexes will be migrated the next time that the STRDSMOND command runs.
- Life of Data and Indexes. The length of time in days to maintain index data on the system. For migration, this value must be greater than the Migrate After *n* Days value.

OnDemand does not migrate index data from the database to archive media until migration processing runs. The STRDSMOND command is the migration utility. You can control automatic migration processing by scheduling the command to run with the appropriate options. You can also manually start migration processing by running the command from the command line.

After a migrated table is successfully loaded into the System Migration application group, the table is deleted from the database. However, OnDemand keeps track of all migrated tables. That way, if index data in a migrated table is needed, then OnDemand can alert an administrator to take action (such as **manually** import the table back into the database).

How to migrate indexes

You can control automatic migration processing by scheduling the STRDSMOND command to run with the appropriate options. You can also manually start migration processing by running the STRDSMOND command from the command line.

When the STRDSMOND command migrates indexes, it saves the following messages in the System Log. A set of three messages should be saved in the System Log for each table that is migrated from the database to archive storage:

```
166 ApplGroup Segment Export (ApplGrp) (Segment)
14 DB Info Exported (SQL Code)
87 ApplGrp Load (System Migration)
```

The first message identifies a table of application group index data that is to be migrated from the database to archive storage. The second message reports the status of exporting the table from the database to temporary storage. The third message reports the loading of information about the migrated table into the System Migration application group. The System Migration application group must be assigned to a storage set that identifies an archive storage media type (such as optical or tape).

Chapter 3. Migrating and importing index data

Index migration is the process by which OnDemand moves index data from the database to archive storage. This process optimizes database storage space while allowing you to maintain index data for a very long time. You typically migrate index data after users no longer need to access the information, but for legal or other business requirements, you still need to maintain the data for some number of years. If a user queries index data that has been migrated, an administrator must import a copy of the migrated table into the database. After maintaining the imported table in the database for a specified number of days, OnDemand deletes it from the database.

This section provides information about importing index data into the database, including what happens when a user queries for migrated data, how to import the index table or tables required by the query, and what happens after you import a table into the database.

We assume that an experienced OnDemand administrator will use the information provided in this section. If you have questions about any of the topics in this section or if you would like help configuring your system to support migrating and importing of index data, please contact the IBM support center.

In general, migrating index data is not recommended.

Configuring the system

There are a number of things you should consider regarding system configuration before you make any decisions about index migration.

System Log messages

OnDemand provides the System Log for administrators to monitor the system. When you install and configure OnDemand, you initialize the System Log tables. The System Log is critical to the operation of the system.

When OnDemand processes a query for application group indexes that have been migrated to archive storage, it saves a message in the System Log and sends a message to the System Log user exit program. A message is also sent to the QSYSOPR message queue.

You can configure the system to examine the messages that OnDemand sends to the System Log user exit and mail them to an administrator or send them to another program. You can also configure the System Log user exit program to determine what action to take when a user queries for data that has been migrated to archive storage.

System Log user exit program

When a client queries index data that has been migrated to archive storage, OnDemand saves message number `168 ApplGrp Segment Not Available` in the System Log. OnDemand also sends the message to the System Log user exit program. If you have defined your own System Log user exit program, then you can determine the action to take when OnDemand sends the message to the System Log user exit program. For example, you may want the program to notify

© Copyright IBM Corp. 2001, 2002

an administrator that a request for a table of migrated index data has occurred. Please see the *IBM Content Manager OnDemand for iSeries Common Server Planning and Installation Guide* for more information about the System Log user exit program.

Archive Storage Manager

Before OnDemand can migrate index data to archive storage, you must configure a migration policy with the information that the archive storage manager uses to maintain the data. The migration policy should maintain the data indefinitely. If you need the system to maintain a backup copy of the index data, then you should specify this in the migration policy.

Storage sets

OnDemand uses the System Migration application group to manage all index data that is migrated to archive storage. You must assign the System Migration application group to a storage set that identifies an archive storage media type.

Application groups

When you define an application group, you specify the storage management information that determines how long OnDemand maintains data stored in the application group and when OnDemand takes certain actions. For example:

- Life of Data and Indexes. Determines the length of time that OnDemand maintains index data and report data stored in the application group.
- Migration of Indexes. Determines the number of days before OnDemand moves index data from the database to archive storage.

  You should plan to migrate index data only after users no longer need to access the reports to which it refers. Only in exceptional situations should users need to access index data that has been migrated. If a user needs to access index data that has been migrated to archive media, the process of importing the table back into the database requires **manual** actions by an administrator, and usually results in a significant delay in completing the query. The import process also requires additional space in the database to hold the imported tables, additional log file storage, and temporary storage on the server to run the import process.
- Keep Imported Migrated Indexes. Determines how long that OnDemand maintains the imported index data in the database before it is scheduled for deletion.

If you need to maintain index data in archive storage, then you must configure the Migration of Indexes in your application groups. You must specify the number of days to keep the index data on disk in Keep Imported Migrated Indexes. OnDemand will schedule imported index data for deletion from the database after it resides in the database for the number of days specified in Keep Imported Migrated Indexes or Life of Data and Indexes, whichever occurs first.

You can use the administrative client to configure your application groups.

What happens when a user queries migrated data

There are several ways to be notified when a user requests queries migrated index data.

Message to the user

When the server determines that the index data required to complete a query has been migrated to archive storage, it sends a message to the client program. The message states that the data required to complete the query is not available and that the user should contact an administrator.

Message to the System Log

When OnDemand determines that the index data required to complete a query has been migrated to archive storage, it saves a message in the System Log. An administrator can open the System Log folder to search for and display messages in the System Log.

OnDemand also sends a message to the QSYSOPR message queue and the System Log user exit program. You can configure the system to examine the message and send an alert to an administrator or call another program to take some action.

Importing index data

If index data is to be imported back from archive media, it is important to consider the following:

Verify database storage space

Importing migrated index data from archive storage back into the database requires additional database storage. Before you import the index data, you should verify that sufficient free space is available.

Verify database log file space

Importing migrated index data from archive storage back into the database requires database journal storage. Before you import the index data, you should verify that sufficient free space is available.

Run the STRIMPOND command

OnDemand provides the STRIMPOND command to import tables of migrated index data from archive storage back into the database. (The name of the application group and the index table to import can be obtained from the message that OnDemand saved in the System Log.)

After the STRIMPOND command completes the import operation, you can open the System Log folder to see the messages that were generated by the import process. The messages will reference the ARSADMIN program name.

See online help for more information about the STRIMPOND command and its parameters.

After you import index data

After you import index data from archive storage back into the database, you should notify the user to retry the query.

Expiring imported migrated indexes

OnDemand schedules an imported index table for deletion after it resides in the database for the number of days specified in the Length of Time to Keep Imported Indexes property in application groups. After an imported index data reaches the

specified value, the next time that the STRDSMOND command runs, the imported index table is deleted from the database. (However, the table still exists in archive storage.)

You typically configure the STRDSMOND command to run automatically on a regular schedule. You can also run the STRDSMOND command manually.

Configuring index migration

If you find that your users are often querying for index data that has been migrated to archive storage, then we recommend that you configure your application groups to increase the length of time that OnDemand maintains the index data in the database. This should reduce the number of queries that need migrated index data.

Keeping imported migrated indexes

OnDemand schedules imported index data for deletion after the index data resides in the database for the number of days specified in Keep Imported Migrated Indexes or Life of Data and Indexes, whichever occurs first.

Part 2. Storage administration

This part is of primary interest to administrators responsible for configuring and maintaining disk and archive storage. This part is organized as follows:

- Document storage
- Backup and recovery

© Copyright IBM Corp. 2001, 2002

Chapter 4. Document storage

This section is organized as follows:

- Defining document storage management
- Migrating documents
- Removing documents

Defining document storage management

The document storage management definitions determine where and when OnDemand stores documents and how it maintains them.

When you load a document into OnDemand, you assign it to an application group. The application group is the last document storage management component that you define, because it requires storage set and migration policy definitions, which you must create first. The application group identifies the storage set and determines where documents should be loaded. You assign each application group to a storage set. You can load documents onto disk, onto archive media, or onto both disk storage and archive storage. The disk storage manager maintains documents on disk. The archive storage manager maintains documents on archive media. The archive storage manager uses a migration policy to determine where to store documents and how long to maintain them. After a document ages for the specified number of days, the migration process can move it from disk to archive storage.

Application groups

The application group is the last component that you must define because it requires storage set and migration policy definitions. The application group provides a way to group related documents. All documents in the application group are loaded on the media that is part of the storage set to which the application group is assigned. All documents in the application group migrate according to the rules that are defined for the application group's migration policy.

Use the administrative client to create the application groups that determine the document storage for your documents. You typically define one application group for each set of your documents that have similar storage requirements. For example, documents that must be retained for a specific length of time, in specific storage locations and stored on specific types of media.

Loading

An application group definition contains the rules for loading documents into an application group. It requires one or more storage sets, which you must create first. The application group determines if documents are loaded onto disk, archive storage, or both. If the application group causes documents to be stored only on disk, then the migration policy specifies when (or if) documents are copied to archive storage.

See the following properties on the Storage Management tab of the application group:

- Storage Set. Determines where documents will be loaded.

© Copyright IBM Corp. 2001, 2002

Note: If you specify Cache Only, then documents can only be loaded onto disk.

- Cache Data. Determines if documents will be loaded into disk storage.

  Note: If the storage set is a cache-only storage set, then documents must be loaded onto disk.

- Migrate Data from Cache. If you specify When Data is Loaded, then documents will be loaded into archive storage. (Migration is disabled for the application group.)

Migrating

Migration is the process of copying documents from disk to archive storage as controlled by the rules of the application group's migration policy. However, because a document is eligible to be migrated does not mean that it will be migrated. Other factors affect migration, such as the frequency with which you run migration processing (migration cannot take place until you run migration processing).

A migration policy contains the rules for migrating the documents in an application group. The Storage Management tab defines how long a document stays on disk and, through the storage set and migration policy, where the document will be moved next. The migration policy level identifies the next location.

See the following settings on the Storage Management tab:

- Storage Set. Determines the next location for documents.

  Note: If you specify Cache Only, then migration is disabled for the application group.

- Migrate Data From Cache. Determines when documents are eligible to be migrated.

  Note: If you specify No or When Data is Loaded, then migration is disabled for the application group.

Disk storage manager

The disk storage manager (initiated using the STRDSMOND command) maintains documents on disk. Documents migrate from disk storage to archive storage based on the migration policy that is defined for the application group. The disk storage manager can delete documents after they exceed the **Cache Data for n Days** or **Life of Data**, whichever occurs first. See "Removing documents" on page 18 for more information.

Archive storage manager

The archive storage manager (initiated using the STRASMOND command) is the interface to the archive media (typically optical or tape). The archive storage manager maintains a backup or long-term copy of documents. Before loading documents, you must define storage sets and migration policies, optical volumes, tape devices, and tape volumes. The archive storage manager can delete documents after they exceed the **Retention Value** specified in the migration policy. See "Removing documents" on page 18 for more information.

Migrating documents

OnDemand provides automatic migration to copy documents from disk storage to archive storage (for documents that were not loaded to archive storage) and to make documents eligible for deletion to maintain free space on disk. Migration helps to ensure that there is sufficient free space on disk, where faster response time can provide the most benefit to your users.

Important:

1. You should run migration processing on a regular schedule to make sure that a backup copy of your documents gets created as soon as practically possible. If you defer the migration of documents to archive storage, and disk storage were to become corrupted, then you could be left without a backup copy of your documents.
2. The STRASMOND command must only be run in batch (SBMJOB parameter set to *YES). Running this command interactively (with SBMJOB(*NO)) may cause SQL errors.
3. The STRASMOND command does not have a specific parameter for specifying an instance name. By default, the QUSROND default instance is used, and will produce the desired results for most systems. If you need to name a different instance, you must specify it as INSTANCENAME/POLICYNAME in the Policy parameter of the STRASMOND command. You do not need to put your entry in apostrophes. If you enter an instancename/policyname, the command prompter will automatically enclose it in apostrophes. This will force the value to be processed in whatever case it was entered (uppercase, lowercase, or mixed case). Note that if you initiate the archive storage manager by running the STRDSMOND command with RUNASM(*YES), then the instance name is passed from the disk storage manager and no further specifications are needed.
4. If you run STRDSMOND for a specific application group (rather than the default of *ALL) and you set the Run ASM (RUNASM) parameter to *YES, be aware that ASM will run for ALL application groups, even though you have named a specific application group for DSM to use. You can, however, name a specific Policy for ASM to process, if desired. Also note that when you specify RUNASM(*YES), OnDemand will initiate a separate batch job for ASM.

You control automatic migration processing by scheduling the STRDSMOND and STRASMOND commands to run with the appropriate options. See your operating system information for details about how to schedule tasks. You can also start migration processing by running the command manually.

The STRDSMOND program uses an application group's storage management information to control when migration for an application group occurs:

- If you use **Next Cache Migration** to control when migration for an application group occurs, then the disk storage manager runs migration processing each time that you start the STRDSMOND command with the appropriate options.
- If you use **After n Days in Cache** to control when migration for an application group occurs, then a document must be stored on disk for at least the specified number of days before it is eligible to be migrated.

The disk space that migrated documents occupy can be reused after expiration processing completes. When you run migration processing, you should also run expiration processing so that the disk storage manager can reclaim the disk storage space occupied by migrated documents.

Migration processing in the system log

When you run the STRDSMOND command, it saves messages about its activities in the system log. The types of messages saved in the system log depend on the options that you specify when you run the STRDSMOND command. The number of messages saved in the system log during a migration process depend on the options that you specify for the STRDSMOND command, the number of application groups and segments of data processed. Table 1 lists the messages you could see in the system log following migration processing.

Table 1. Messages from the STRDSMOND command in the OnDemand System Log

Message Number	Message Information	Explanation
110 Cache Migration	(Date) (Server)	About to begin disk migration on the specified server.
197 Cache Migration	(ApplGrp) (ObjName) (Server)	One of these messages for each storage object migrated to archive storage. Migration copies a storage object if its ″After n Days in Cache″ period has passed or the application group uses the ″Next Cache Migration″ migration method.
124 Filesystem Statistics	(filesystem) (% full) (server)	One of these messages each time you run the STRDSMOND command. Message is Informational only, to report the percentage of space used in the system auxiliary storage pool (ASP).

Removing documents

Documents *expire* (are eligible for removal) because their disk expiration date or archive retention period has passed. Expired documents can then be removed by the storage managers. The disk storage manager identifies documents for removal by using the application group's expiration information. The archive storage manager marks documents for removal based on the criteria defined in the migration policy.

Documents expire from disk when they reach their disk expiration date. If a document's disk expiration date is less than its Life of Data period, then the document is simply removed from disk storage. Subsequent requests for the document are satisfied by the archive storage manager. When the document reaches its Life of Data period, information about it is removed from the OnDemand database (the document can no longer be retrieved). When the document's archive retention period has passed, information about it is removed from the archive storage manager database.

Because a document is eligible to be removed does not mean that it will be deleted from storage. The disk storage manager does not delete expired documents from storage until expiration processing runs. During expiration processing, the archive storage manager deletes information about expired documents from its database. However, the actual documents may remain on archive media until such time that the media on which they reside is reinitialized.

Important: The disk storage manager and the archive storage manager delete documents independently of each other. Each uses its own criteria to determine when documents expire and should be removed from the system. Each uses its own utilities to remove documents. However, for final removal of documents from the system, you should specify the

same criteria to the disk storage manager and the archive storage manager. In other words, the sum of the Duration values for all migration policy storage levels plus the value for the Days in Cache should equal the Life of Data and Indexes value in the application group.

Removing documents from disk storage

The expiration policy determines when documents are eligible for deletion from disk. You define the expiration policy when you create the application group. The following properties on the Storage Management tab comprise the expiration policy:

- Cache Data for n Days. The length of time in days to keep documents on disk. After a document reaches this value, it is eligible to be deleted from disk.
- Life of Data. The length of time in days to maintain documents on the system.

  Note: If you specify Never Expire, then expiration processing is disabled for the application group.
- Expiration Type. Determines whether one or more documents are eligible to be deleted at a time. For example, by default, the **Load** expiration type means that the set of documents that were loaded together will also expire together.

  Note: Segment (in contrast to Load) is another possible choice for expiration type. This is the first time that we've mentioned a *segment*. Up to now, we've talked about *documents*, which is the data object that most people associate with the OnDemand system. However, administrators who maintain the system may also work with segments, which represent many documents, and *storage objects*, which are containers of compressed documents that are maintained by the storage managers.

The disk storage manager does not delete expired documents from disk until expiration processing runs. The STRDSMOND command is the expiration utility. You can schedule the STRDSMOND command to run automatically or you can run it manually. You should make sure that the STRDSMOND command runs periodically so that the disk storage manager can reclaim the space that is occupied by expired documents.

Expiration processing in the system log

When you run the STRDSMOND command, it saves messages about its activities in the system log. The types of messages saved in the system log depend on the options that you specify when you run the STRDSMOND command. The number of messages saved in the system log each time that expiration processing runs depends on the options that you specify for the STRDSMOND command and the number of application groups and segments of data processed. Table 2 lists the messages you could see in the system log following expiration processing.

Table 2. Messages from the STRDSMOND command in the OnDemand System Log

Message Number	Message Information	Explanation
109 Cache Expiration	(Date) (Min%) (Max%) (Server)	About to begin disk expiration processing on the specified server.

Table 2. Messages from the STRDSMOND command in the OnDemand System Log (continued)

Message Number	Message Information	Explanation
196 Cache Migration	(ApplGrp) (ObjName) (Server)	One of these messages for each storage object deleted from disk. A storage object is eligible to be deleted when its ″Cache Data for n Days″ or ″Life of Data″ period has passed, whichever occurs first.
124 Filesystem Statistics	(filesystem) (% full) (Server)	One of these messages each time you run the STRDSMOND command. Message is Informational only, to report the percentage of space used in the system auxiliary storage pool (ASP).

Important: In addition to the messages listed in Table 2 on page 19, you should monitor the system log every day for messages that indicate that your disk space is becoming full. The STRDSMOND command automatically saves a message in the system log when the amount of space used in a disk storage file system exceeds a threshold. The threshold is 90%.

Removing documents from archive storage

Note: Removing a document from archive storage means that the backup or long-term copy of the document will be deleted from the system. You typically remove documents from archive storage when you no longer have a business or legal requirement to keep them.

A migration policy specifies the criteria that makes a document eligible for deletion. Documents become eligible for deletion under the following conditions:

- Administrators delete documents from archive media
- An archived document exceeds the time criteria in the migration policy (how long archived copies are kept)

The archive storage manager does not delete information about expired documents from its database until expiration processing runs. You can run expiration processing either automatically or manually using the STRASMOND command (or you can use the STRDSMOND command with the Run ASM parameter set to *YES). You should make sure that expiration processing runs periodically to allow the archive storage manager to reuse storage space that is occupied by expired documents. When expiration processing runs, the archive storage manager deletes documents from its database. The storage space that these documents occupy then may become reusable.

Chapter 5. Backup and recovery

It is important to implement a backup plan that will serve to protect your data archives in the event your organization experiences a disaster, large or small. You need to have a plan and prepare so that you can recover.

Backup considerations

Items you need to consider when planning your regular backups include:

- How to recover the optical or tape media itself

  You should consider having a copy of the media stored at an offsite location.
- How to recover the data on disk

  You should have procedures in place that backup all of the OnDemand data that normally resides on disk, as well as any stored data that has not yet migrated to optical or tape.

Remember: Disk backups will contain database files that reflect the actual archive data location and status at the time of backup. DO NOT restore from a backup that is one week old. Archived data might have migrated to optical or tape, but database control files identify incorrectly the archived data location as disk. In order to minimize this situation, you should save your database library(s) and save your OnDemand Integrated File System directories on the same schedule to keep them synchronized. Perform the backups **at least** as frequently as you run the disk storage management (STRDSMOND) or archive storage management (STRASMOND) commands. (Possibly even more frequently if you often manually delete reports that are stored in OnDemand.)

The following objects need to be saved by the appropriate OS/400 commands:

- The OnDemand licensed program (**5722-RD1**)
- User profiles **QRDARS400, QRDARS4001, QRDARS4002, QRDARS4003, QRDARS4004, QRDARS4005, QONDADM, and QRDARSADM**
- **Authorization list QRDARS400**
- **Database files** in **QUSRRDARS** and **QUSROND** libraries, and any other data libraries that contain data for OnDemand instances other than the default QUSROND instance
- **Output queues** that are being monitored

  These output queues can contain OnDemand processed spooled files. Saving the **output queues** does not save the data in them. There is no save command to save spooled files.
- **Integrated File System directories**

  Each instance that is created in OnDemand has an **Integrated File System directory** that is named the same as the instance. To backup only the OnDemand data that resides on disk, you could backup all items that are found in the **Integrated File System directory /QIBM/UserData/OnDemand/<instance>** where <instance> is the name of your OnDemand instance. Remember that the name of the default instance is QUSROND.
- **Configuration files**

© Copyright IBM Corp. 2001, 2002

See Chapter 9 of the *IBM Content Manager OnDemand for iSeries Common Server Planning and Installation Guide* for more information and a list of configuration files.

Recovery considerations

Contact your OnDemand support provider for instructions on recovering your OnDemand archives in the event of a disaster. Many factors can influence the recovery plan, depending on the frequency and extent of the backups you have available.

Reports

OnDemand can store copies of reports and resources on disk storage and archive storage:

- The primary purpose of disk storage is short-term, high-speed storage and retrieval of reports.
- The primary purpose of archive storage is long-term storage and retrieval of reports. The reports in archive storage can also be used as backup copies in the event that disk storage becomes corrupted or unavailable. Archive storage consists of optical or tape storage volumes.

Most customers configure the system to copy reports to disk storage and archive media at the same time, when they load a report into the system.

OnDemand can retrieve a copy of a report from archive storage after the report has been deleted from disk storage or if the copy on disk storage is unavailable. However, you must configure the system to maintain multiple copies of reports. You configure OnDemand to use archive storage by defining migration policies and storage sets that identify archive storage, assigning application groups to the storage sets, and configuring other storage management parameters in application groups.

Note: If you do not plan to copy reports to archive storage, then we recommend that you take regular backups of the data on disk. However, if a media failure occurs or disk storage becomes corrupted, users cannot retrieve reports until the backups are restored.

Part 3. Using iSeries Navigator for OnDemand administration

This part of the book contains information about using iSeries Navigator to define and maintain objects on the system. For example, you can use iSeries Navigator to define reports to OnDemand and maintain OnDemand users, groups, printers, storage sets and migration policies.

This part of the book describes a typical system, including requirements about users of the system and data to be loaded on the system, and contains examples and ideas about how to implement the requirements. It is organized in roughly the order that you define your first report to OnDemand. It starts with adding users and groups to the system, shows how to define a printer, and then moves into procedures for adding a migration policy and storage set. Later chapters provide detailed instructions about defining a report: adding an application group, an application, and a folder.

© Copyright IBM Corp. 2001, 2002

Chapter 6. Installation

This section describes the hardware, software, memory, and disk space requirements for installing and using the administrative client.

Hardware

The administrative client requires the following hardware:

- An IBM-compatible PC with an Intel Pentium® 166 MHz or faster processor
- A super-VGA display and adapter with a minimum resolution of 800x600
- Physical connection to the network, such as a Token Ring or Ethernet network adapter
- A CD-ROM drive for installation (optional, if you plan to install the administrative client from a network file server)

Software

To install or use the administrative client, you must be running Windows 2000, Windows 98, Windows NT 4.0 with SP5 or later, or Windows XP on your computer. To install the administrative client from a CD-ROM, the CD-ROM drivers must be installed on your computer. To install the administrative client from a network file server, the appropriate network software must be running on your computer. To connect to an OnDemand server, the standard TCP/IP support for the Windows operating system must be running on your computer.

To use the iSeries Navigator interface to the OnDemand Administrator and for additional OnDemand administrative functions, you must install the iSeries Navigator OnDemand Administrative Functions component within iSeries Access Express.

Memory

The administrative client requires a PC with at least 64 MB of memory (RAM). Your PC may need more memory if you run more than one application at the same time.

Disk space

To install OnDemand clients, disk space requirements depend on the features that you select.

You need at least 55 MB of free disk space to install all of the OnDemand features. To use less than 55 MB of disk space, use the Installation Components dialog box in Setup to select only the features you need. To install only the software required for the administrative client, you need at least 12 MB of available disk space.

For more information on disk space requirements for each OnDemand feature, start Setup from the CD-ROM, or follow the Setup steps below:

1. Insert the CD-ROM in the appropriate drive.
2. From the Windows taskbar, click Start, and then choose Run.
3. Type x:\client\win32\setup (where x is the letter of your CD-ROM drive).

© Copyright IBM Corp. 2001, 2002

4. On the Setup Type and Working Directory dialog box, select Custom.
5. On the Installation Components dialog box you can select each component individually and then click Change to see all the specific features and their disk space requirements.

Installing the iSeries Navigator interface to OnDemand

To use the iSeries Navigator interface, the OnDemand Administrative Functions component must be installed. From the Environment tasks window of iSeries Navigator, click on Install plug-ins. This will cause the Plug-in selection window to appear, allowing you to select the plug-ins you want to install. Select to install the OnDemand Administrative Functions plug-in.

For further detailed information on the installation of iSeries Navigator or plug-ins, see the *iSeries Access for Windows – Setup*, SC41-5507-03.

Installing the administrative client

The Setup program transfers the OnDemand client files from the CD-ROM to each user's PC.

Important: If the CD-ROM that you receive contains Client software at a level below 7.1.0.4, you must update the Administrator Client software from the Internet to the 7.1.0.4 level or higher. To download the OnDemand Clients, go to the internet FTP web site at ftp://service.software.ibm.com/software/ondemand/fixes/v71/

You can install all of the OnDemand features at once, or individual features as you need them.

Running Setup

When you run the Setup program, the Setup screens show the names of the OnDemand directories so you know where the files are being placed.

Information about network installations is available in the *IBM Content Manager OnDemand: Windows Client Customization Guide*.

To install on a user's PC

To install from the CD-ROM or to install from a network file server, follow these instructions:

1. Insert the CD-ROM in the appropriate drive or obtain the drive letter of the network drive on which the OnDemand software is located on the network file server.
2. From the Windows taskbar, click Start, and then choose Run.
3. Type `x:\client\win32\setup` (where `x` is the letter of your CD-ROM drive or the network drive).
4. Click OK.

After the Setup program starts, click Next to continue.

On the Setup Type and Working Directory dialog box, select Custom and then click Next to continue.

On the Installation Components dialog box, select Administrator. (You can also select other components that you want to install on the PC.)

Click Next to continue. Follow the instructions on the screen to complete the installation.

To install on a network file server

Information about network installations is available in the *IBM Content Manager OnDemand: Windows Client Customization Guide*.

To use automated install

Automated install allows administrators to standardize the OnDemand installation for all users in an organization. It also allows administrators to install OnDemand clients without the presence of users at their PCs. To read more about automated install, see the *IBM Content Manager OnDemand: Windows Client Customization Guide*.

To uninstall

You can use the Add/Remove Programs applet in Control Panel to remove the OnDemand client from your PC. The Add/Remove Programs applet removes the directories for OnDemand and any references to OnDemand in system files.

To run the uninstall

1. From the Windows taskbar, click Start. Choose Settings, and then choose Control Panel.
2. Double-click Add/Remove Programs.
3. From the list, select OnDemand32.
4. Click Add/Remove.
5. Click OK.

Chapter 7. Introduction

The administrative functions in OnDemand are set up using iSeries Navigator and the OnDemand Administrator, and include:

- Defining reports to the system
- Adding and maintaining OnDemand users and groups
- Adding and maintaining server printers
- Maintaining storage sets and migration policies
- Generating reports about users, groups, applications, application groups, storage sets, folders, and printers
- Adding and maintaining servers
- Setting system parameters for OnDemand servers and client programs
- Copying items from one OnDemand server to another
- Tracking changes made to the system. When you use the iSeries Navigator to add or update the database, OnDemand places a document in the system log which shows the changes that you made.

The following is a list of the administrative items that are maintained directly through iSeries Navigator:

- Tape devices
- Tape volumes
- Optical storage groups
- Optical volumes
- Disk pool storage groups
- Monitor definitions
- Migration policies

The items that are maintained through the OnDemand Administrator are:

- Users
- Groups
- Applications
- Application groups
- Storage sets
- Folders
- Printers

The OnDemand Administrator is launched by right-clicking on Common Server Administration within the OnDemand Archive section of iSeries Navigator. From the resulting pop-up menu, click Client Administrative Functions to launch the OnDemand Administrator.

Getting started

- You can collapse and expand areas in the navigator pane (on the left) to make it easier to see. A plus sign next to an area means there are items inside.
- When you click on an area in the navigator pane, the items appear in the list pane (on the right).

© Copyright IBM Corp. 2001, 2002

- To make the panes narrower or wider, point to the vertical bar between the two panes of the window until the pointer turns into a two-headed arrow. Then click and hold the left mouse button and drag it in either direction.
- Use buttons on the toolbar to switch between the different ways to look at items: large or small icons, a list, or details.
- To query the server for a new list of items, press the F5 key or select Refresh List from the View menu.

Using online help

Online help provides information to assist you with completing tasks. The administrative client online help contains information about the options, fields, and commands on the windows, dialog boxes, and property sheets that you see when using the program.

To display the online help, press F1 any time the administrative client is active in Windows. Help is available for dialog box commands and options. The main help topic for each dialog box usually contains information about the kinds of tasks you can perform. For example, the online help about Logical Views lets you learn how to create public and private logical views. The online help provides brief procedures rather than lengthy descriptions.

Adding a server

You can use the New Server command to add a server.

1. In the OnDemand section of iSeries Navigator, right click Common Server Administration. In the resulting context menu, click Client Administrative Functions, then logon.
2. Click once on OnDemand Servers at the top of the list of servers in the left panel of the OnDemand Administrator.
3. From the File menu, select New Server to open the Add a Server dialog box.
4. Type the name of the server in the Server field. The server name identifies the server in the navigator pane of the administrator window. You can use an alias, the actual computer or network name of the server, or any other identifier you choose. By default, the administrative client copies what you type to the Host Name field.
5. Verify the value of the Host Name field. (By default, the Host Name field contains the same value as the Server field.) The host name can be a host name alias, fully-qualified host name, or IP address of the server.
6. Verify the Protocol. Choose from TCP/IP or Local:
   - TCP/IP. Use TCP/IP (Transmission Control Protocol/Internet Protocol) as the network protocol. To use TCP/IP, the server and the client must include TCP/IP in the protocol stack.
   - Local. Use the local file system manager to communicate with the server. This is most often used when your archived data is on a CD-ROM.
7. For TCP/IP, verify the Port Number that the library server monitors for client program requests.

   The default value, `0` (zero), means that the server monitors the port number of the OnDemand TCP/IP service. By default, the server monitors port number 1445. If you plan to use a port number other than 1445, then you must enter a valid port number. If you have more than one instance defined to OnDemand, then you would have separate server definitions for every instance with unique port numbers for each instance.

8. You should clear the Attempt Unified Logon check box. (Unified Logon is used only for Windows servers.)
9. If you are adding a Local server, then enter the full path name of the Directory in which the local database will be created in the space provided. Use the Browse command to select a directory.
10. Click OK to add the server.

Logging on a server

1. Point to the server and double click the left mouse button to open the Logon dialog box.
2. Enter your OnDemand userid and password in the spaces provided and click OK.

Changing passwords

Note: This section applies only if you are using OnDemand userids and passwords rather than OS/400 userids and passwords. See the *IBM Content Manager OnDemand for iSeries Common Server Planning and Installation Guide* section titled *OnDemand user relationship to OS/400 user profiles* for important information about passwords.

Changing a password

To change a user's password:

1. Select and expand the library server.
2. Select Users.
3. In the User ID list, point to the userid and click the right mouse button.
4. From the pop-up menu, select Update to open the Update a User dialog box.
5. Type the new password in the Password field.

   Note: When creating a password, the value that you specify can be a maximum of 20 characters. However, the password authentication that is built into OnDemand verifies only the first eight characters that are entered by the user. The additional characters are provided for customers that choose to implement their own password security by using the logon user exit. Contact the IBM support center for more information about the logon user exit.
6. Verify the new password by retyping it into the Verify Password field.
7. Click OK. OnDemand updates the database and returns to the main window.

Changing an expired password

OnDemand provides password expiration processing to help you manage security on the system. You can set a value that represents the time in days that passwords assigned to users remain valid. After a user's password reaches the value that you specify, the user must change the password. See "System parameters" on page 32 for information about how OnDemand expires passwords.

After a password reaches the expiration value, the next time the userid is used to log on to a server, OnDemand prompts the user to enter a new password.

The user must enter the current password for the userid, a new password, and verify the new password by retyping the new password

System parameters

OnDemand system parameters allow you to establish the following operational settings for client programs and servers.

Maximum Password Age

Sets a time limit for passwords and determines when OnDemand prompts users to change passwords. The default setting is `Password Never Expires`, meaning that passwords do not expire and OnDemand never prompts users to change passwords. If you select `Password Always Expires`, then users must change to new passwords each time that they log on to a server. To set a specific time limit for passwords, select `Expires In __ Days` and enter the number of days that passwords are valid in the space provided. The value can be from `1` (one) to `365`.

If you are linking your OnDemand userids to OS/400 user profiles, then Maximum Password Age should be set to Password Never Expires. If you specify a value for Maximum Password Age, then OnDemand may force a user to change their password before it is required by OS/400.

Minimum Password Length

Determines whether passwords are required. If passwords are required, also determines the fewest number of characters that passwords can contain. The default value is `At Least 8 Characters`.

If you select `Permit Blank Password`, meaning that passwords are not required, then the valid password length is 0 (zero) to 20 characters.

If you are linking your OnDemand userids to OS/400 user profiles, then Minimum Password Length should be set to Allow Blank Password. This prevents OnDemand from trying to impose its own rules on the length of a password and allows OS/400 to use its own rules.

To set a specific minimum password length, select `At Least __ Characters` and enter a number in the space provided. The value can be from 1 (one) to 20 characters.

Note: When creating a password, the value that you specify can be a maximum of 20 characters. However, the password authentication that is built into OnDemand verifies only the first eight characters that are entered by the user. The additional characters are provided for customers that choose to implement their own password security by using the logon user exit. Contact the IBM support center for more information about the logon user exit.

Inactivity Time Out

Determines when OnDemand terminates sessions between inactive clients and the server. The default setting is `Time Out in 60 Minutes`. The value can be from `1` (one) to `1440` (24 hours).

The period of inactivity is measured between requests to a server. For example, when a user enters a query, OnDemand searches the database and builds the document list. This completes a request to the server. If the user does not work

with the items in the document list, open another folder, or invoke another query before the inactivity time out occurs, OnDemand automatically terminates the session with the client.

Use caution when you set the inactivity time out. For example, assume that you set the inactivity time out to 10 (ten). You log on to OnDemand to add an application group. Creating the application group takes you 15 minutes to complete. After entering all of the information about the application group, you click OK to create the application group. OnDemand issues a message that shows a time out has occurred. You must logoff the server, and you cannot save the information you entered about the application group.

System Logging

Determines the messages that OnDemand saves in the system log. OnDemand provides the system log to help you track activity and monitor the system. OnDemand saves messages that are generated by the various commands, such as the ADDRPTOND command. OnDemand can save a message in the system log when the following events occur:

- A user logs on to the system
- A user logs off the system
- A user logon fails
- Application group data is queried, retrieved, loaded, updated, deleted, or maintained

User Exit Logging

OnDemand provides a user exit at each of the four system log event points. These exits allow you to filter the messages that OnDemand saves in the system log and take action when a particular event occurs. For example, you can provide a user exit program that sends a message to a security administrator when someone attempts and fails to log on to the system. You can also use a user exit to determine what information appears in the system log.

If you plan to migrate index data to archive storage, then we recommend that you configure the system to save application group messages in the system log and send them to a system log user exit program. You should design a system log user exit program to notify an administrator when a query for migrated data occurs. Before a query for migrated data can be completed, an administrator must import a copy of the table or tables that are required from archive storage to the database.

See the *IBM Content Manager OnDemand for iSeries Common Server: Planning and Installation Guide* for help with configuring the system log user exit.

Login Processing

The login processing system parameters allow you to specify whether userids and passwords are case sensitive.

Before continuing with this section, please refer to the OnDemand userid relationship to OS/400 user profiles section in the *IBM Content Manager OnDemand for iSeries Common Server: Planning and Installation Guide* for a detailed explanation of the significance of choosing to relate your OnDemand users to existing OS/400 user profiles. It is important that you understand that concept before you make your choices for Login processing.

Now that you understand the relationship between OnDemand users and your OS/400 users, you should note the following:

- If your OnDemand userids **are** linked to your OS/400 user profiles (which is the default when OnDemand is installed), and if your OS/400 security level is set to 0 or 1, you **SHOULD NOT** check the Password Case Sensitive checkbox.
- If your OnDemand userids **are** linked to your OS/400 user profiles (which is the default), and if your OS/400 security level is set to 2 or 3, you **SHOULD** check the Password Case Sensitive checkbox.
- Regardless of your system security level, you should **NEVER** check the UserID Case Sensitive checkbox if your OnDemand users are linked to your OS/400 user profiles.
- If you are using OnDemand userids and passwords that are **not** linked to OS/400 user profiles (which is not the default), then you can set the two Login processing checkboxes as you choose.

By default, userids and passwords are case insensitive. When you add a user, OnDemand converts lowercase letters in the userid to uppercase. A person can type letters in a userid in uppercase, lowercase, or mixed case letters. For example, if you add the userid `LaGuarde`, a person can enter `LAGUARDE`, `laguarde`, or `LaGuarde` to log on to the server.

If you select UserID Case Sensitive, then a user must type the userid exactly as it was entered when the user was added. For example, if you add the userid `LaGuarde`, then the user must enter `LaGuarde` to log on to the server.

If you select Password Case Sensitive, then a user must type the password exactly as it was entered when the user was added. For example, if you set the password to `Spring2Far`, then the user must enter `Spring2Far` to log on to the system.

We strongly encourage you to decide whether you want userids and passwords to be case sensitive when you install the system, change the defaults if necessary, and do not change the settings again. Otherwise:

- If userids are initially case insensitive and you later choose UserID Case Sensitive, then userids that were added before you changed the parameter must be entered in uppercase. The same is true for passwords.
- If userids are initially case sensitive and you later clear UserID Case Sensitive, then the userids that were added before you changed the parameter that contain mixed or lowercase letters will no longer be valid. The same is true for passwords.

Setting system parameters

To set the system parameters for an OnDemand server:

1. Log on to the server.
2. Point to the server and click the right mouse button.
3. From the pop-up menu, select System Parameters to open the System Parameters dialog box. Figure 1 on page 35 shows an example of the System Parameters dialog box.

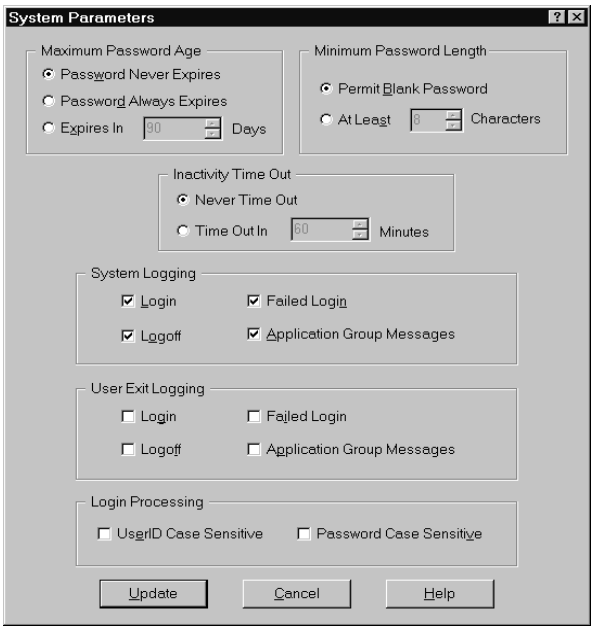


Figure 1. System Parameters

4. To change the Minimum Password Age, select the appropriate option. If you select `Expires In`, enter the number of days in the space provided.
5. To change the Minimum Password Length, select the appropriate option. If you select `At Least`, enter the number of characters in the space provided.
6. To change the Inactivity Time Out, select the appropriate option. If you select `Time Out In`, enter the number of minutes in the space provided.
7. To choose a System Logging, User Exit Logging, or Login Processing option, select the check box next to the item.
8. When you have finished making changes to the system parameters, click Update. OnDemand stores the changes in the database and returns to the administrator window. (Click Cancel to close the System Parameters dialog box without saving your changes.)

Adding items to a server

This section explains how to add items to a server. You can use commands or a drag-and-drop operation to add items to a server.

Note: When you use the administrative client to add or update the database, OnDemand adds a record to the system log that shows the changes you made.

New command

After logging on to a server, select the area, for example, Users. From the File menu, select the New command to open the Add dialog box.

Copy command

After logging on to a server, select the area. In the list pane, point to the item that you want to copy and click the right mouse button. From the pop-up menu, select Copy to open the Add dialog box. The fields in the Add dialog box contain

information copied from the item you selected. Before you can add the item, you must change the item name. Depending on the item you want to add, you may need to change other fields.

Export command

The Export command is like the Copy command, except OnDemand adds the item to a different server. You can use the Export command to export items from one server (the source server) and add them to another (the destination server).

After logging on to the server that contains the item you want to export, select the area. In the list pane, point to the item that you want to export and click the right mouse button. From the pop-up menu, select Export to open the Export dialog box. Verify the destination server. Then click Export to add the item to the server.

Note: If the item exists on the destination server, the export fails.

Drag and drop operation

You can copy and export items using a drag-and-drop operation. For example, to export items from one server and add them to another, select one or more items from the list pane and, while holding the left mouse button down, point to the destination server. Then release the mouse button. If you are logged on to the destination server, then OnDemand opens the Export dialog box. If you are not logged on to the destination server, then OnDemand opens the Logon dialog box. After verifying options in the Export dialog box, click Export to copy the items to the server.

Note: If the item exists on the destination server, the export fails.

You can also use a drag-and-drop operation to copy an item. For example, to copy a user, select the userid from the User ID list and, while holding the left mouse button down, point to the same server on which the user is listed. Then release the mouse button to open the Add dialog box.

Report Wizard

OnDemand provides user assistance and easy-to-use tools to help you administer OnDemand. The Report Wizard helps you add a report to OnDemand by asking questions, which allows you to progress in an organized manner toward completing an application group, application, and folder. Here are a few things to remember about the Report Wizard:

- You move through the Report Wizard by answering questions that appear on the screen.
- You can return to the previous screen at any time by clicking Previous.
- You can advance to the next screen at any time by clicking Next.
- You can advance to the final screen by clicking Finish. By choosing Finish, you permit the Report Wizard to make all remaining decisions for you.
- You can obtain online help for a screen at any time by clicking Help or pressing F1.

Important: The Report Wizard processes your own input files. You must select a sample input file to proceed.

You can use the Report Wizard to add an application group, application, and folder for a selected report. These actions include defining indexing information,

defining database and folder fields, configuring data and storage management, specifying whether the application group can contain more than one application, and naming the application group, application, and folder.

You can also use the Report Wizard to add an application to an existing application group. This action includes defining indexing information, specifying storage information, and identifying the application within the application group. To add an application to an application group, the application group must have a database field to hold the values that uniquely identify an application within the application group. The field must contain at least one unassigned application identifier. See the Field Information page in application groups for detailed information about application identifiers.

Starting the Report Wizard

From the administrative client, log on to the server to which you want to add the report.

- To define a new application group, application, and folder, click the Report Wizard icon on the toolbar.
- To add an application to an existing application group:
  - Under the server, select Application Groups
  - Select the name of the application group to which you want to add the application
  - Click the Report Wizard icon on the toolbar

Follow the on-screen instructions to add the report.

Using the Report Wizard

The screens in the Report Wizard are described below. On most screens, standard options will already be selected for you. Unless you have a clear reason not to, accept the defaults.

Note: Depending on how you use the Report Wizard, you may not see all of the screens described below.

Introduction screen

Provides a brief explanation of the Report Wizard (see Figure 2 on page 38. First choose the data type of the report you are defining. Click Select Sample Data button to select a file that contains a sample of the actual report data. The Wizard lets you select a spooled file on the server (search by user profile or by output queue) and copies that sample data to your workstation for you to use for indexing.

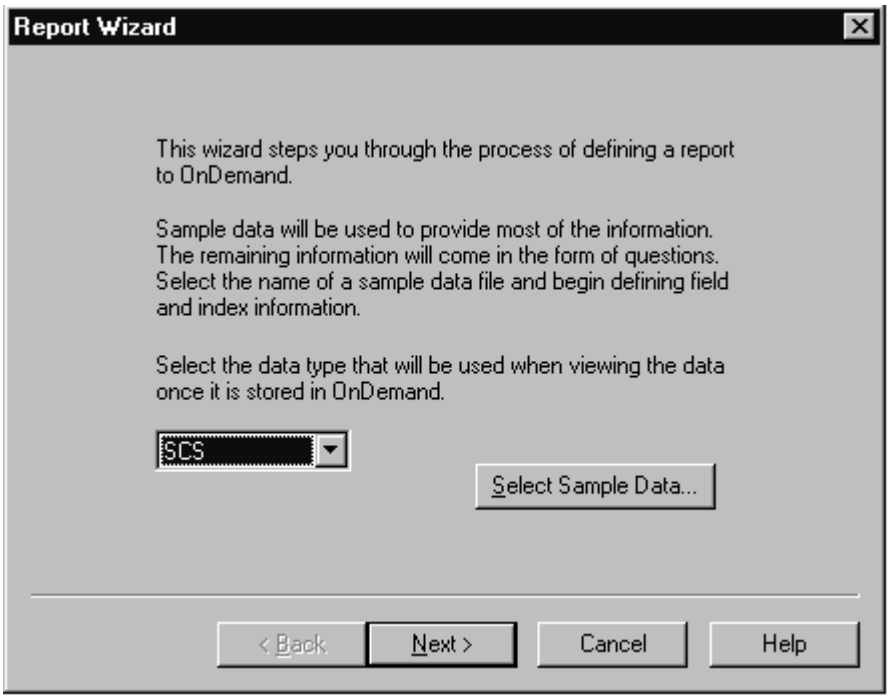


Figure 2. Report Wizard

When you click OK, the Report Wizard reads the data into the Report window.

Report window

Displays the sample data file and provides easy-to-use tools to help you define indexing information, database fields, and folder fields. Press F1 to display the online help for options and commands available from the Report window. Use the online help to learn how to define triggers, fields, and indexes, database fields, and folder fields.

Important: When you have finished defining the indexing, database, and folder information, be sure to save your changes when prompted.

Managing data screen

When you load a report into the system, you can specify that you want report data to be stored in Large Objects. You also need to specify how you want OnDemand to manage annotations that users attach to pages of the report.

Application identifier screen

When you use the Report Wizard to add an application to an existing application group, you must specify the name of the application and select a value that uniquely identifies the application within the application group.

Storage management screen

Determines where the storage manager maintains copies of reports, and how and when OnDemand deletes report data from the system.

Applications in the application group screen

If the report that you are defining is one of several that will be stored in the same application group, then you can use the Report Wizard to define the database field (called an application ID field) that contains the values that identify an application within the application group, define the folder field that users use to search a specific application, and specify the length of the application ID field. It is highly

recommended that you ALWAYS define an Application ID Field on the Field Information page. (Figure 22 on page 78 shows where you would define one.) You may not think you need one at the time you create your application group, but for flexibility in the future, you may need one later. The important thing to note is that an application ID field CANNOT be added to the application group AFTER it is created; it must be part of the original application group definition. Also note that the application ID field can be hidden so that end-users are not confused by it if it does not provide value to them. Refer to online help for more information about the Application ID Field.

Name screen

Specify the names of the application group, application, and folder. After you enter the names, OnDemand queries the server to make sure that the names are valid and unique.

Wizard complete screen

Confirms the selections you made for the report. Click Display to view details about the application group, application, and folder. From the detail report window, choose the Print icon from the toolbar to print a copy of the detail report.

Note: When you are satisfied with the selections you made for the report, click Finish to complete defining the report. OnDemand adds the application group, application, and folder to the library server, closes the Report Wizard, and returns to the administrator window.

Chapter 8. Concepts

You can use iSeries Navigator to setup and maintain OnDemand tape devices, tape volumes, optical storage groups, optical volumes, disk pool storage groups, output queue monitors, and migration policies.

You can use the OnDemand Administrator within iSeries Navigator to maintain OnDemand users, groups, printers, storage groups, application groups, applications, and folders. To use this tool, right click on Common Server Administration within the OnDemand Archive section of iSeries Navigator. Select Client Administrative Functions from the resulting menu. This will launch the OnDemand Administrator.

Migration Policies

Migration policies contain migration and storage media characteristics for data archived using OnDemand. The information is used by the archive storage management process (ASM), also referred to as the migration process, to determine if and when archived data should be moved as it ages through a hierarchy of storage media having different performance and capacity charactersitics. Examples of these media types include disk, optical, and tape storage. Each step in the movement of data through this storage hierarchy is referred to as a migration policy storage level, or simply, a storage level. Each policy must contain at least one storage level. Additional levels may be defined to meet your storage and retrieval requirements.

The migration policy also specifies:

- if separate archived files are to be aggregated, or combined, with other archived files having similar retention and migration characteristics which can increase the performance of the migration process by reducing the total number of files to be processed
- if two copies of archived data are to be kept at some or all levels in the migration sequence
- if a one-time backup is performed

The migration policy name can be up to 60 characters long and must not be a duplicate of another policy within the same instance. If the Enable aggregation is selected, the archive storage management process combines individual archived objects on disk into larger objects to provide efficient processing. This process occurs prior to migration of the object from disk to the first storage level. The aggregation process appends to the same file (aggregate) until the aggregate is 'closed.' The aggregate is closed when it either reaches a specified maximum size or a specified number of days. Storage levels can be added before or after a pre-existing storage level. Existing storage levels can also be changed or removed. If a one-time tape backup is requested, a tape media type must be specified.

Monitor Definitions

Monitor definitions are used to specify what output queue will be monitored for spooled files to be processed. If defining more than one monitor job, specify a unique job name for each monitor. OnDemand will only process spooled files that are in a ready (RDY) state. When the monitor job selects a spooled file from the

© Copyright IBM Corp. 2001, 2002

selected output queue for archiving, it needs to determine which application group and application to associate with the spooled file so that the file can be archived correctly. Since the only data available to the monitor are the attributes of the selected spooled file, the application group name and application name must be derived from the contents of one of these attributes. Not all attributes are suitable for this purpose. OnDemand will examine the contents of up to three of the following nine attributes, in the order specified in the Check first, Check next, Check last selections:

- Spooled file name
- Form type
- User data
- Job name
- User-defined options 1 through 4
- User-defined data

The attribute selected from the 'Check first' pulldown list is examined first. If the value of this attribute does not match the name of an existing application group, OnDemand examines the attribute selected from the 'Check next' pulldown list, if specified. If the value of this attribute does not match an application group name, the attribute selected from the 'Check last' pulldown list, if specified, is checked. If a valid application group is not determined using the above method, the spooled file is moved to another output queue designated as an 'error' queue which is defined in this monitor definition. If the spooled file is successfully archived, it can be moved to a 'processed' queue or deleted according to the specifications in the monitor definition. The same process is followed to find a valid application name, unless you specify that the application name is the same as the application group name.

The monitor can be started manually, by a job scheduler, or started when the subsystem starts. The monitor can be ended manually, after a specified time period, after all queue entries are processed, or it can be specified when the monitor is started.

Tape devices

When a tape backup is requested, or if you will use tape as an archive media, you must defined a tape device to OnDemand. The tape device name you specify must match the name of an existing OS/400 tape device description. A media type must be specified from the list of supported media types for either read or write operations. If the tape device being defined has an automatic cartridge loader (ACL), the number of cartridges can be specified. A value of zero specifies that this device does not have an ACL. OnDemand can be used with an automated tape library; if a tape library will be used, the name must be specified. If a tape manager other than OnDemand (such as BRMS) is used, a media library does not need to be specified.

Tape volumes

Tape volumes that can be used by OnDemand must be defined. The name of the volume specified must match the name that was used when the tape volume was initialized. The capacity and media type of the volume must be specified. The media type tells OnDemand which tape device to use. A media device library can be specified if an automated tape library is used. Leave this field blank if a tape manager other than OnDemand (such as BRMS) is used. A tape volume can be

marked full, preventing OnDemand from writing any additional data to the volume. OnDemand will automatically mark the volume full when it detects that the tape is full.

Optical storage groups

Optical storage groups are used by OnDemand to group sets of optical volumes together to store related data. A storage group lets you group together reports that have similar storage requirements such as days on disk or expiration days. By referring to a specific storage group in your migration policy, you can control which reports are stored on a particular set of optical volumes. If optical volumes are defined as rewritable, the space can be reused on the volume by defining a free space threshold percent and volume full reset. If the volume full reset is not defined, once the volume is marked full it remains full unless manually changed. The optical storage group can also be defined as the primary or backup storage group.

Disk pool storage groups

The disk pool storage group is used to identify an OS/400 auxillary storage pool (ASP) that the archive storage management process may use as storage media when migrating archived data. You must specify a pool number ranging from 1 to 32 which corresponds to an existing ASP. The type of data, primary or backup, that will be stored to the defined ASP must also be specified.

Users

When you define an OnDemand user, you create a userid with which a person in your organization logs on to the OnDemand server. You can optionally add the userid to folders and application groups permissions, which is one way to let the user open folders and access data.

Each person in your organization logs on to the server using an OnDemand userid. OnDemand authenticates the userid and determines the usage and administrative authority available to that person, based on the userid. It is important that you understand the details in the OnDemand userid relationship to OS/400 user profiles topic in the *IBM Content Manager OnDemand for iSeries Common Server Planning and Installation Guide* (SC27–1158) before you continue with this section.

An OnDemand userid does not necessarily have to identify an individual by name. However, for accounting purposes, most customers assign an OnDemand userid to each person that will use the system. When you initialize the system, OnDemand automatically creates the QONDADM userid. The QONDADM userid has system administrator authority. A system administrator can perform the basic user functions, such as logging on the system and opening folders, and administrative functions, such as defining users and groups to OnDemand, and creating, updating, and deleting application groups, applications, folders, migration policies and storage sets, and printers.

Note: The QONDADM userid has an initial password of QONDADM1. Because the QONDADM userid has system authority, we force you to change the password the first time you sign on.

When naming OnDemand users, the name that you specify:

- Can contain one to 10 or 128 characters (bytes) depending on whether you are linking OS/400 user profiles to your OnDemand users. If you are linking the two, then the user names should match your OS/400 user profile names.
- Cannot include the ' (apostrophe), * (asterisk), % (percent), + (plus), _ (underscore), [(left bracket),] (right bracket), " (double quote), or blank character
- Must be unique to the server

You can specify a userid in mixed case. By default, OnDemand ignores the case (for example, `LaGuarde` is the same as `laguarde`). OnDemand converts lowercase letters in a user name to uppercase (`LaGuarde` is stored as `LAGUARDE`). However, depending on how you configure the Login Processing system parameters, userid processing on your system may be different (the case may be significant). See "System parameters" on page 32 for more information.

User types

When you add a user to OnDemand, you specify the *User Type*. The User Type and the Authority determines the types of tasks that the user can do when logged on to the system. You can choose from the following User Types:

User Users can log on to OnDemand, open folders that they are authorized to access, and search for and retrieve data from application groups that they are authorized to access. Users can be given authority to do other things on the system.

User Administrator
A user that can also add, update, and delete users and user administrators. A user administrator can be given authority to do other things on the system.

Application Group/Folder Administrator
A user that can also add, update, and delete application groups, applications, and folders. An application group/folder administrator is automatically given Logical Views permission to all application groups. An application group/folder administrator can be given authority to do other things on the system.

System Administrator
A user that can also add, update, and delete any user, group, application group, application, or folder on the system. A system administrator is automatically given Logical Views permission to all application groups. A system administrator is also the only user that can maintain storage sets and printers.

Note: When adding or updating a user, you are not permitted to set the User Type or Authority to a level that exceeds your own. For example, a user with Create Users and Create Groups authority cannot create a user with Create Folders authority.

By default, only the user, the user that created the user, user administrators, and system administrators can view or maintain the user. See the User Permissions page for more information.

Only a system administrator, an application group/folder administrator, a user with administrator authority for an application group, or a user with add document permission can store data in an application group.

Only a system administrator, an application group/folder administrator, or a user with delete document permission can delete data from an application group.

Users who need to run server commands such as those listed in Appendix A, "Command reference" on page 133 or server APIs from QSHELL such as those listed in Appendix D, "API and user exit reference" on page 145 need to have QRDARSADM as the group profile (or a supplemental group) in their OS/400 user profile.

Authority

The Authority options allow the user to do other things in OnDemand. For example, A User Type of User can be permitted to create users. The authority options that you can select depend on the User Type. For example, if the User Type is Application Group/Folder Administrator, then by definition, the user can create application groups and folders. Therefore, the only additional authorities that the user can be given are Create Users and Create Groups. Choose from the following:

Create Users
An optional authority for users and application group/folder administrators:

- If the User Type is User, lets the user create users with a User Type of User.
- If the User Type is Application Group/Folder Administrator, lets the user create users with a User Type of User or a User Type of Application Group/Folder Administrator.

Note: Users with Create Users authority can maintain the users that they create, so long as they remain an administrator of the user.

Create Groups
An optional authority for users, user administrators, and application group/folder administrators. Lets the user create groups. Users with Create Groups authority can maintain the groups they create, so long as they remain a group owner.

Note: These group definitions are not the same as OS/400 group profiles, although the names may match if you find that easier to maintain.

Create Application Groups
An optional authority for users and user administrators. Lets the user create application groups. Users with Create Application Groups authority can maintain the application groups that they create, so long as they remain an application group administrator.

Create Folders
An optional authority for users and user administrators. Lets the user create folders. Users with Create Folders authority can maintain the folders that they create, so long as they remain a folder administrator.

Note: When adding or updating a user, you are not permitted to set the User Type or Authority to a level that exceeds your own. For example, a user with Create Users and Create Groups authority cannot create a user with Create Folders authority.

Groups

OnDemand groups are a means to organize users of the system by function, authorization, or any other purpose that you might require. You do not have to assign a user to a group, however doing so can simplify administration of users with similar requirements and capabilities.

When you define a group, you can add users to the group and specify folder and application group permissions that are common to all of the users that belong to the group. The permissions determine the types of actions users assigned to the group can perform on the system.

When naming groups, the name that you specify:

- Can contain one to 20 characters (bytes)
- Cannot include the ' (apostrophe), * (asterisk), % (percent), + (plus), _ (underscore), [(left bracket),] (right bracket), " (double quote), or blank character
- Can be mixed case. However, the case does not create a unique name (`LaGuarde` is the same as `laguarde`)
- Must be unique to the server

You can assign a group owner. The group owner can add users to and remove users from the group. To maintain a group's application group and folder permissions, the group owner must have administrator authority for the application groups and folders or be an application group/folder administrator or a system administrator. If you do not assign a group owner, only a system administrator user can maintain the group.

Note: These group definitions are not the same as OS/400 group profiles, although the names may match if you find that easier to maintain.

Printers

OnDemand supports two types of server print devices: a fax machine and a physical printer. A server print device always has an output queue on the server and is defined using the OnDemand iSeries Navigator administrative client.

PSF/400 is required for formatting Advanced Function Presentation™ (AFP™) print output from OnDemand. (PSF/400 may also be required by some fax software as well.) PSF/400 allows you to use electronic forms, images, graphics, and typographical fonts in the documents that you print. PSF/400 accepts input data streams, such as AFP and line data and prepares the data for the destination print device.

When a user selects an item and chooses server print, OnDemand retrieves a copy of the item and places it on the output queue associated with the server printer.

Storage sets

Storage sets are defined for application groups with similar storage management characteristics, such as the length of time that files are maintained in the application group and the type of media on which the files are stored. A storage set is created automatically when you create a migration policy, and the names are identical. You do not have to manually create a storage set.

If you plan to maintain application group data in archive storage, then you must specify a storage set name in the application group definition that matches the migration policy to be used. Also, the sum of the Duration values for all migration policy storage levels plus the value for the Days in Cache should equal the Life of Data and Indexes value in the application group.

Application groups

An application group represents the data that you store in OnDemand and the documents that users query, view, print, and fax using OnDemand client programs. For example, the data can be reports generated by an application program, index data, and annotations created by users.

When you define an application group, you specify properties of the application group, such as the organization of the database and the storage characteristics for the files that are to be stored in the application group. You also define the database fields that will hold index data extracted from the reports that you store in the application group.

OnDemand extracts index data from the reports that you load into an application group and places the data in the database fields that you define. OnDemand uses the index values to identify the documents that meet the search criteria entered by a user.

When you define an application group, you can also select the types of application group messages that OnDemand saves in the system log.

When you define an application group, you specify permissions that let users access and maintain the application group and application group data. You can identify the groups and users that can access data stored in the application group with OnDemand client programs. You can specify the types of functions that users can perform, such as viewing, printing, and annotating reports. You can assign administrator authority to a user or a group. Administrator authority allows a user to update the application group, for example, to authorize other users to access data stored in the application group.

OnDemand organizes information about an application group into tabs: General, Message Logging, Storage Management, Permissions, Field Definition, and Field Information. Each tab contains options, fields for you to enter information about the application group, and command buttons.

Applications

You typically define an application for each different type of report or source of data that you plan to store in OnDemand.

When you define an application, you assign the application to an application group and specify the physical and logical characteristics of the report. The physical characteristics of a report include the code page, the type of data found in the input file, and information about carriage control characters. The logical characteristics of a report include the different ways that you want to present the information contained in the report to your users.

The OnDemand data indexing, loading, and viewing programs use the information that you provide to process the report. For example, you can specify the parameters that the OnDemand indexing program uses to locate and extract index

data from the report. You can create logical views for the application. Logical views represent different ways to display pages of the report. You can set up printing options, such as defining a default printer for users and printing options for AFP and line data documents.

OnDemand organizes information about an application into tabs: General, View Information, Indexer Information, Load Information, Logical View Fields, Logical Views, and Print Options. Each tab contains options, fields for you to enter information about the application, and command buttons.

Folders

A folder provides users the means to access the reports that you store in OnDemand. A user opens a folder, constructs a query, and retrieves documents from the application groups that can be searched from the folder. The user can use the folder to view, print, annotate, fax, and e-mail documents.

When you define a folder, you specify the properties of the folder, such as the name and description of the folder, create the search and display fields that appear when the user opens the folder, and map the folder fields to application group database fields.

You can also specify the groups and users that can open the folder with OnDemand client programs and other folder permissions. For example, you can authorize a user to be the folder administrator. The folder administrator can authorize other users to open the folder and make changes to the folder fields.

OnDemand organizes information about a folder into tabs: General, Permissions, Field Definition, Field Information, and Field Mapping. Each tab contains options, fields for you to enter information about the folder, and command buttons.

About application groups, applications, and folders

Before you can store a report into OnDemand, you must create an application group and an application. Before users can search for and retrieve data, you must create a folder.

- Users open a folder to search for reports that you load into OnDemand. You define one or more search fields for the folder. A folder search field is mapped to an application group database field.
- Each database field that you define represents a category of information in the report, such as a customer name, invoice number, or balance. When you add an application group, OnDemand creates a database table. The database fields that you define are columns in the table.
- You define an application for each report that you plan to store in OnDemand. When you add an application to the system, you define information that the OnDemand viewing, indexing, and loading programs use. When you define indexing information, you identify the name, location, and length of each category of index information that you want to extract from the report. When you define loading information, you map index fields in the report to application group database fields.
- When you load a report into the system, OnDemand stores the index values extracted from the report into database fields in records that are added to an application group table.
- OnDemand uses the index values to identify the documents that meet the search criteria that the user entered into the folder search fields.

When you want to define a report to OnDemand, your first task is to identify the application group from which OnDemand obtains information about the index fields and how documents are to be maintained on the system. When you define an application group, you specify how you want OnDemand to structure information in the database and define the database fields. When you define an application group, you also specify how you want OnDemand to maintain data on the system. For example, you might specify that report data should be maintained on disk storage for 60 days and in archive storage for five years. OnDemand maintains all of the data stored in the application group the same way. OnDemand maintains each report that you store in the application group for the same length of time.

- You can store the report in an existing application group. However, you must be able to index the report using the database fields that are already defined in the application group. You cannot add or remove database fields from an application group. The storage management information for the application group must support the length of time that you want OnDemand to maintain the report on the system and how and where that you want OnDemand to store and maintain the report data.

  You can verify information about an existing application group with the Properties command. The General tab shows the database organization for the application group. The Storage Management tab shows the data migration information. The Field Definition tab shows the application group database fields.
- If there are no application groups defined to OnDemand or there are no application groups that support the database and storage management requirements of the report, then you must add an application group to the system.

After you add an application group, you must define an application for the report. Most customers create an OnDemand application for each different type of report or source of data that they plan to store in OnDemand. When you create an application, you must assign it to an application group. The application group determines where OnDemand will store the report data. When you create an application, you also specify information that the OnDemand client programs use to view and print pages of the report and you specify instructions for the data indexing and loading programs.

The last step in the process of adding a report to OnDemand is to create a folder. Users open the folder to search for, display, and print reports. When you define a folder, you select the application group that contains the data that you want users to search when they open the folder. By creating folders that can search specific application groups, you can determine the reports that are available to users when they open a folder.

When you define a folder, you define search and display fields. You specify characteristics of the folder fields, such as default search operators and whether OnDemand displays default search values for the fields when a user opens the folder. You also map the folder fields to application group database fields.

OnDemand permissions

Permissions are the means by which OnDemand determines who can open folders and search for documents stored in application groups. OnDemand also uses permissions to determine who can maintain folders and application groups with the administrative client.

By default, only an application group/folder administrator, a system administrator, or the person who adds the folder can open and maintain the folder. By default, only an application group/folder administrator, a system administrator or the person who adds the application group can access data stored in the application group or maintain the application group.

OnDemand provides several ways for you to specify permissions. You can specify permissions when you add or update a folder or an application group. You can also add, remove, or update a user's or group's folder or application group permissions when you add or update the user or group.

As both a convenience and security measure, you can assign a user to a group. OnDemand groups allow you to organize users by function, authorization, or any other logical grouping that you might require. When you assign a user to a group, the user obtains the permissions that are in effect for the group. For example, suppose that you create a group and authorize the group to open the Student Information folder. Any user that you assign to the group automatically obtains permission to open the Student Information folder. If you assign a user to more than one group, the user normally obtains the permissions of all of the groups. However, there are exceptions. See "Specifying permissions" on page 51 for details.

Note: These group definitions are not the same as OS/400 group profiles, although the names may match if you find that easier to maintain.

You can specify a default set of permissions for folders and application groups. OnDemand uses the default permissions when users and groups do not have specific permissions for the folder or application group. If you specify permissions for a group, then the group permissions take precedence over the default permissions. If you specify permissions for a user, then the user permissions take precedence, regardless of any group that the user may belong to or the default permissions that you specified.

Folder permissions

You can specify default (*PUBLIC) folder permissions and folder permissions for specific groups and users. The default permissions provide every user and group defined to the server with the permissions that you specify. Permissions for a group provide the users that you add to the group with the permissions that you specify for the group. Permissions for a user provide the user with the permissions that you specify. By default, only an application group/folder administrator, a system administrator or the person who adds the folder can open and maintain the folder.

Refer to online help for details about setting folder permissions.

Application group permissions

You can specify default (*PUBLIC) application group permissions and application group permissions for specific groups and users. The default permissions provide every user and group defined to the server with the permissions that you specify. Permissions for a group provide the users that you add to the group with the permissions that you specify for the group. Permissions for a user provide the user with the permissions that you specify. By default, only an application group/folder administrator, a system administrator or the person who adds the application group can access data stored in the application group and maintain the application group.

Refer to online help for details about setting application group permissions.

Specifying permissions

To ease the administration of OnDemand, most customers organize their users into groups, add the groups to folders and application groups, and specify permissions for the groups. You should plan your groups before you begin creating them. After you start using the system, you may find it difficult to change the organization of your groups.

Note: These group definitions are not the same as OS/400 group profiles, although the names may match if you find that easier to maintain.

When you add a user to a group, the user automatically obtains the permissions that were specified for the group. When you add a user to more than one group, the user normally obtains the permissions of all of the groups. For example, using the group properties listed in Table 3, a user that belongs to both groups can open the Student Bills and Student Transcripts folders.

Table 3. Group permissions

Group	GID	Folders	Permission
Admissions	80100	Student Transcripts	Access
Accounting	80101	Student Bills	Access

Most situations involve adding a group to a folder, specifying permissions for the group, and then adding users to the group. However, there may be situations when you need to deny a group of users access to a folder. When you use groups to deny access to a folder, you must understand how OnDemand determines folder permissions for a group (and users assigned to the group). For example, consider the group properties listed in Table 4.

Table 4. Group permissions

Group	GID	Folders	Permission
Admissions	80100	Student Bills	None
Admissions	80100	Student Transcripts	Access
Accounting	80101	Student Bills	Access
Accounting	80101	Student Transcripts	Access

A user that belongs to both groups can not open the Student Bills folder.

When a user belongs to more than one group, OnDemand uses the group identifier (GID) to determine the user's permissions. When two (or more) groups provide permissions for the same folder, the user obtains the permissions of the group with the lowest GID.

It is important to note that you cannot change a GID after a group has been created. You can, however, modify the value that is presented as the suggested GID when the group is created.

In the example depicted in Table 4, both groups have been added to the Student Bills folder. Since the Admissions group has a lower GID than the Accounting group, OnDemand uses the permissions specified for the Admissions group to determine the permissions of a user that is assigned to both groups. Consequently, users assigned to both groups cannot access the Student Bills folder.

You're probably asking yourself, "Why would I assign a user to more than one group" or "Why would I create a group with no access to a folder?" Perhaps some examples will help answer these questions (and clarify the OnDemand permission hierarchy). As you review the examples, please remember the following rules:

- By default, only an application group/folder administrator, a system administrator, or the person who created the folder can access the folder
- You can use the *PUBLIC name to specify default permissions for all other users
- You can specify permissions for specific groups and users:
  - All of the users that belong to a group that you add to a folder will obtain the permissions that you specify for the group
  - A user that belongs to two (or more) groups that have been added to the same folder will obtain the permissions of the group that has the lowest GID
  - The permissions that you specify for a user override all other permissions, including any default permissions (*PUBLIC) and any groups to which the user belongs and that are added to the folder

Examples

The examples that follow show how to add groups to folders and specify folder permissions. The same considerations hold true for adding groups to application groups and specifying application group permissions.

Providing a group of users access to a folder: Let's say that you want to provide a single group of users access to a folder. Complete the following steps:

1. With *PUBLIC selected, clear all of the permissions check boxes (this is the default).
2. Add the group to the folder.
3. Select the Access check box.

Users assigned to the group automatically obtain permission to open the folder.

Denying a group of users access to a folder: Let's say that you want to prohibit a single group of users from accessing a folder, while allowing all other users defined to the server to open the folder. Complete the following steps:

1. With *PUBLIC selected, select the Access check box (this lets all users open the folder).
2. Add the group to the folder.
3. Clear all the permissions check boxes.

Users assigned to the group cannot open the folder.

Providing one group of users access and denying another group of users access to the same folder: Let's say that you want to allow a group of users to access a folder. However, you need to prohibit certain users in the group from accessing the folder. You could exclude the users from the group that can access the folder. However, there may be other folders that you want the users to access as part of the group. To solve this, create two groups, one without access to the folder and the other with access to the folder, and assign the users to the respective groups. For example:

1. Create the "no access" group. This group must have a lower GID than the "access" group. Add users to the group.
2. Create the "access" group. Add users to the group.
3. With *PUBLIC selected, clear all of the permissions check boxes (this is the default).

4. Add the “access” group to the folder.
5. Select the Access check box.
6. Add the “no access” group to the folder.
7. Clear all of the permissions check boxes.

If you later need to deny other users access to the folder, simply add the users to the “no access” group. You can also move users from one group to the other.

Denying one user in a group access to a folder: Let’s say that you want to prohibit one user in a group from accessing a folder. After adding the group to the folder and specifying the access permission, all users assigned to the group can open the folder. To override the group permissions, we can add an individual user to the folder and set permissions at the user level. Complete the following steps.

1. With *PUBLIC selected, clear all of the permissions check boxes (this is the default).
2. Add the group to the folder.
3. Select the Access check box.
4. Add the user to the folder.
5. Clear all the permissions check boxes.

Even though the user belongs to the group, the user cannot open the folder.

Providing one user in a group administrator authority: Let’s say that you want to provide one user in a group the ability to administer the folder. Complete the following steps.

1. With *PUBLIC selected, clear all of the permissions check boxes (this is the default).
2. Add the group to the folder.
3. Select the Access check box.
4. Add the user to the folder.
5. Select the Administrator check box.

Only the user that you added can administer the folder; the other users in the group can open the folder.

Specifying default permissions: The default (*PUBLIC) permissions that you specify for an application group or a folder will apply to every user or group defined to the server who is not provided with specific permissions.

For example, suppose that you specify Access as the default permission for an application group. Every user and group that is not provided with specific permissions can access the data that is stored in the application group. Then, you specify Access as the default permission for a folder. Every user and group that is not provided with specific permissions can open the folder. Later, you add a user, without specifying application group or folder permissions. The user can open the folder and access the data stored in the application group.

While default permissions do provide flexibility to maintain your system, you must make sure that using the default permissions does not circumvent your security strategy. Rather than specifying default permissions for application groups and folders, you may want to use groups as a means to implement your security strategy. For example, you can clear all of the permissions under *PUBLIC and then add groups to a folder and specify the appropriate permissions for each

group. When you add a user to the system, you can assign the user to a group. The user automatically obtains the permissions of the group. If the group does not have access to a particular application group or folder, then neither does the user. With this strategy, until you assign the user to a group, or provide the user with specific permissions, the user cannot access the folder.

Hints and tips

1. To simplify the task of providing access to application groups and folders, give access to a group rather than a user. When a new user needs access, add the user to the group.
2. To allow an Application Group/Folder Administrator to see groups in the permissions list, add the Application Group/Folder Administrator to the groups that require access to application groups and folders.
3. To allow multiple users to administer the same groups, create a group of users and make that group the group owner for any groups that need to be administered by multiple users.
4. The Create Groups authority is most effectively used if it is combined with the Create Users authority or added to a User Administrator. Because the purpose of a group is to give a set of users permissions to another object, it is not very useful if the user that creates the group does not have access to any users. Otherwise, the user that creates a group must be given access to each user that needs to be added to the group.

Chapter 9. Examples

The examples that follow demonstrate how to use the OnDemand iSeries Navigator component to:

- Create tape devices that are available for use by OnDemand if tape archival is selected.
- Add tape volumes which are available for use by OnDemand
- Create optical storage groups which are used to group sets of optical volumes together to store related data.
- Add optical volumes which are available for use by OnDemand.
- Create a disk pool storage group definition to identify an OS/400 Auxiliary Storage Pool (ASP) that the archive storage management (ASM) process may use as storage media when migrating archived data.
- Create monitor definitions which contain information about the output queue to be monitored, the handling of the spooled files being archived, and information about starting and ending the monitor job.
- Create migration policies which provide migration and storage media characteristics for data archived using OnDemand.

The examples that follow also demonstrate how to use the OnDemand Administrator to:

- Add users and groups to the system. Users on the system obtain permission to open folders and access application group data from a group.
- Define a server printer that can be used as the default printer for an application.
- Define a migration policy and storage set.
- Add a report to the system. To add a report, we'll create an application group, an application, and a folder.

System configuration

As you'll read below, we will include an external customer (Customer XYZ) in our examples. Even though we're allowing customers to access the system, for purposes of demonstrating how the system works and how to use the administrative client to implement our requirements, these users could easily represent another department within our company.

We're making the following assumptions about the system:

- Two groups of users: the customer service department and users at Customer XYZ.
- Add the telephone bill report to the system. The telephone report is generated by an application program running on an iSeries system with OnDemand installed.
- Disk storage. When we load a report into the system, OnDemand stores a copy on disk and maintains the report on disk for 60 days.
- Archive storage. When we load a report into the system, OnDemand stores a copy of the report in archive storage. The archive storage manager maintains the report on optical storage for five years.

© Copyright IBM Corp. 2001, 2002

Adding tape devices

1. To add a tape device to OnDemand, left click on the plus sign in front of Common Server Administration, right click on Tape Devices and select New Tape Device.
2. First, the name of the tape device you want to add to OnDemand must be specified. The name must match the name of an existing OS/400 tape device description. In our example, the name of our tape device is TAP01. You may also provide a brief description for your tape device.
3. From the list of supported media types, select the media type for this tape device. In our example, the media type is 8MM.
4. Once you select the media type, the Add button for the Read and Write operations is no longer grayed out. Click on the add button to add the 8MM media type to the Read and Write operations. If the device you are using has an automatic cartridge loader, you can specify the number of cartridges the loader holds. In our example, we will leave this at 0. We will also leave the media library blank. If you are using a media library such as the IBM 3494 tape library with OnDemand, specify the library name here.
5. Click on OK to create the tape device. See Figure 3.

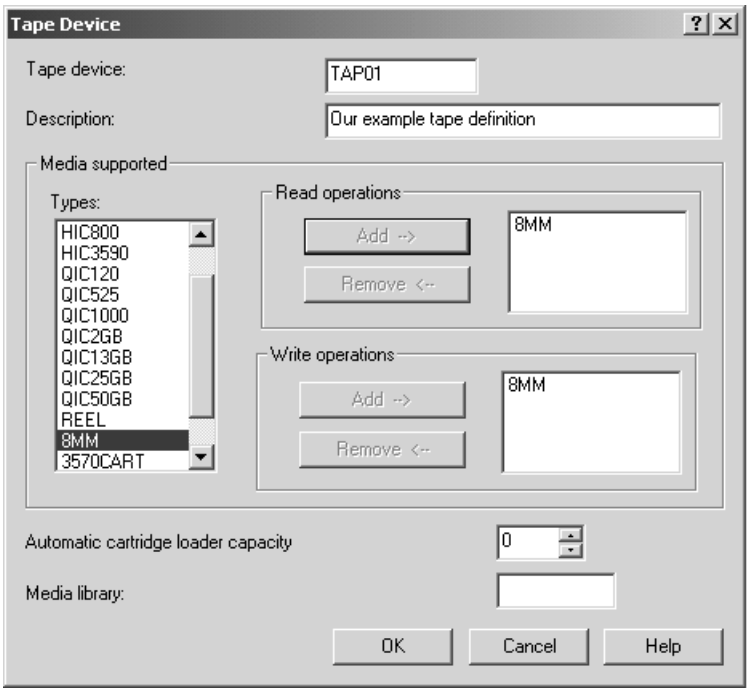


Figure 3. Tape device

Adding tape volumes

1. To add tape volumes to OnDemand, left click on the plus sign in front of Common Server Administration, right click on Tape Volumes and select New Tape Volume.
2. First, we need to specify the name of the tape volume we want to make available to OnDemand. The volume name must match the name you used when you initialized the tape volume on your iSeries server. For this example, the name of our volume is OND001.

3. Next, specify the total capacity of the tape volume in megabytes. In our example, we will specify our capacity of 2300 MB.
4. For media type, use the pull down to select the correct type. For our example, we will select 8MM.
5. For this example we will leave media device library blank. Also, make sure Tape is full is not checked.
6. In our migration policy, which we will create later, we are going to request a tape backup. So, in this example, we need to make our volume type Backup. Then click OK to create our volume. See Figure 4.

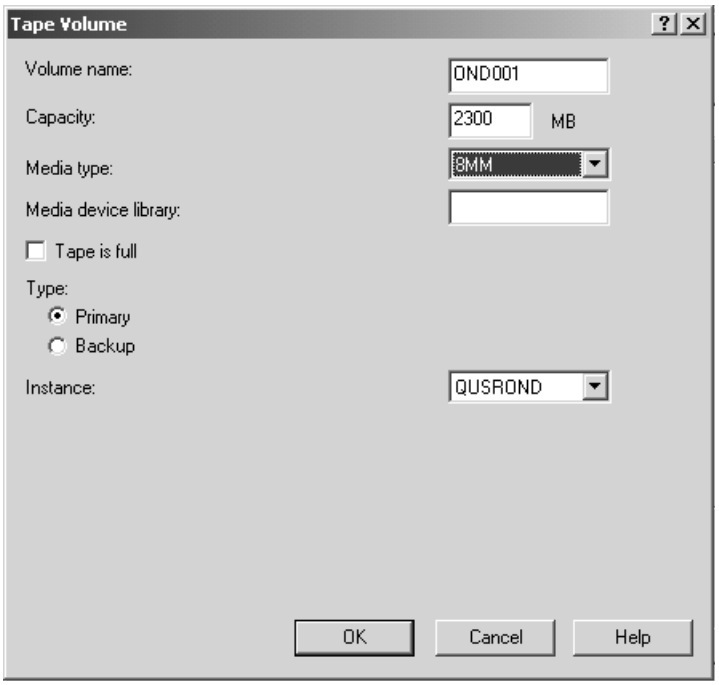


Figure 4. Tape volume

Creating optical storage groups

1. To create an optical storage group, left click on the plus sign in front of Common Server Administration, right click on Optical Storage Groups and select New Storage Group.
2. The optical storage group is used to group sets of optical volumes together for use by OnDemand. For our example, we will call our optical storage group OPTSTG. You may also provide a brief description.
3. Our optical volumes are rewritable, so for this example, we will check the Volume full reset box and set our Free space threshold percent at 40 by clicking on the up arrow. This indicates when the volume full flag is reset. In our example, when there is 40 percent free space available on the volume, the full flag is reset and the volume can again be used for newly migrated data.
4. This optical storage group will contain primary volumes so we will mark the type as Primary.
5. We will leave the instance as the default QUSROND and click OK to create the optical storage group. See Figure 5 on page 58.

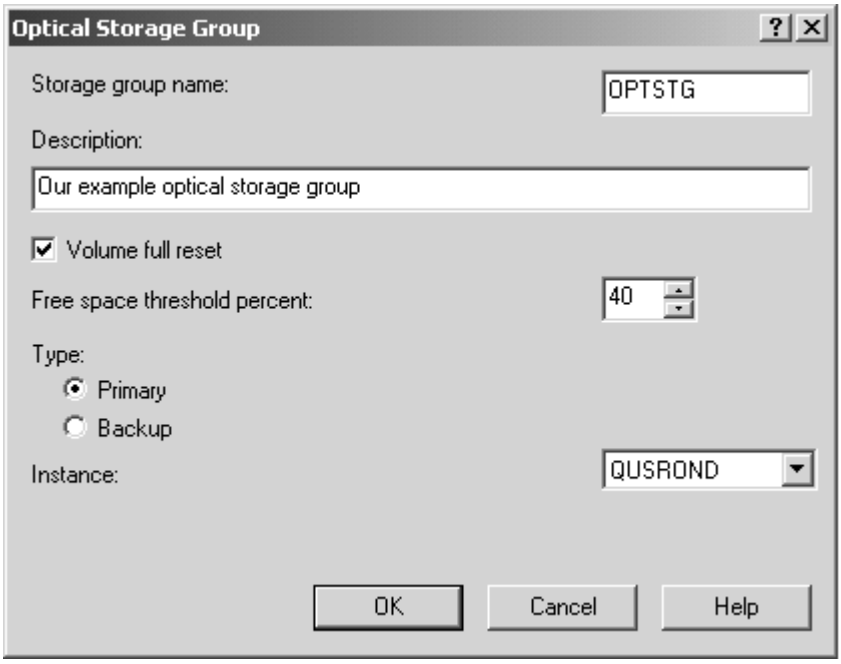


Figure 5. Optical storage group

Adding optical volumes

1. To add optical volumes to OnDemand, left click on the plus sign in front of Common Server Administration, right click on Optical Volumes and select New Common Server Optical Volume.
2. The volume name that we specify must match the name used when the volume was initialized in the optical library. A single disk has two sides which are considered volumes. As you will see in a later step, we must specify the opposite side volume name. The name we will use for this volume is OND00001.
3. Remember that we said our optical storage group was Primary, which means that all the volumes we assign to that storage group must also be of type Primary. Therefore, we will mark this volume as Primary.
4. We will again leave the instance as the default of QUSROND.
5. For this example, we will use a 2.6 GB optical disk. The capacity is for one side of the disk (one volume), so for our example we should specify 1300 MB.
6. The volume we are defining is a rewritable volume. To select the Optical media family, use the pull down and select REWT. We can also use the pull down for the Optical storage group and select the storage group we created earlier, OPTSTG. Our volume is defined as REWT, so it is not necessary to specify the Optical library. (See the online help for more information.) You may specify the library for documentation purposes if you like. In our example, we will leave this blank.
7. We will leave the volume marked as not full. This flag will automatically be set to full when the optical volume has reached its capacity. Remember when we created our storage group, we set the free space threshold to 40%. When this volume has 40% free space available, the full flag will be unchecked and OnDemand can again write to this volume.
8. We will specify the name of the opposite side of this volume as OND00002. Click OK and this will create our volumes. See Figure 6 on page 59.

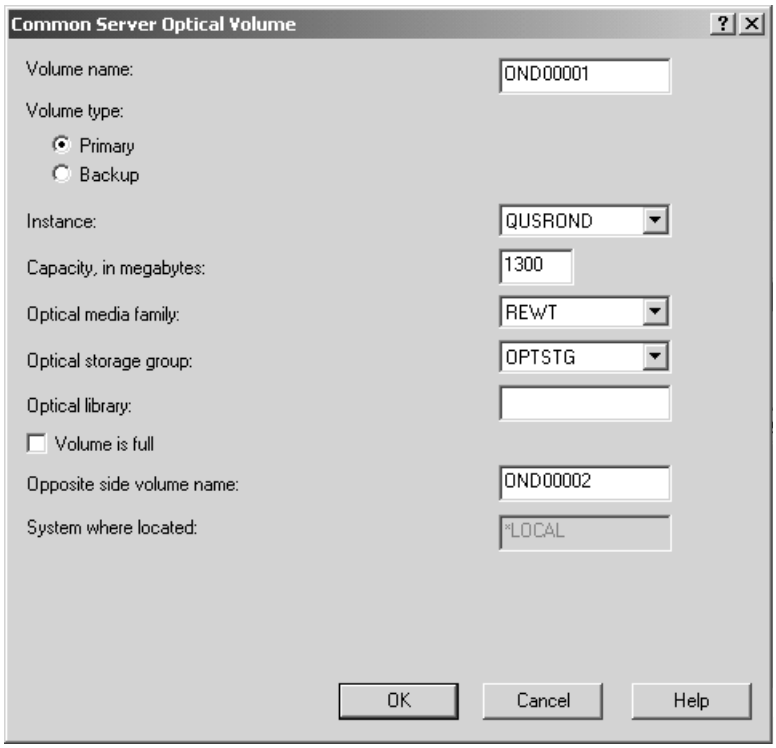


Figure 6. Optical volume

Creating disk pool storage groups

1. To create a Disk pool storage group, left click on the plus sign in front of Common Server Administration, right click on Disk Pool Storage Groups and select New OnDemand Disk Pool.
2. We have chosen to identify an OS/400 Auxiliary Storage Pool (ASP) that the archive storage management (ASM) process may use as storage media when migrating archived data. For our example, we will specify the ASP as 3. You may also provide a brief description for the disk pool storage group.
3. The type for our disk pool is Primary and we will leave the instance as the default QUSROND. Click OK and this will create our disk pool storage group. See Figure 7 on page 60

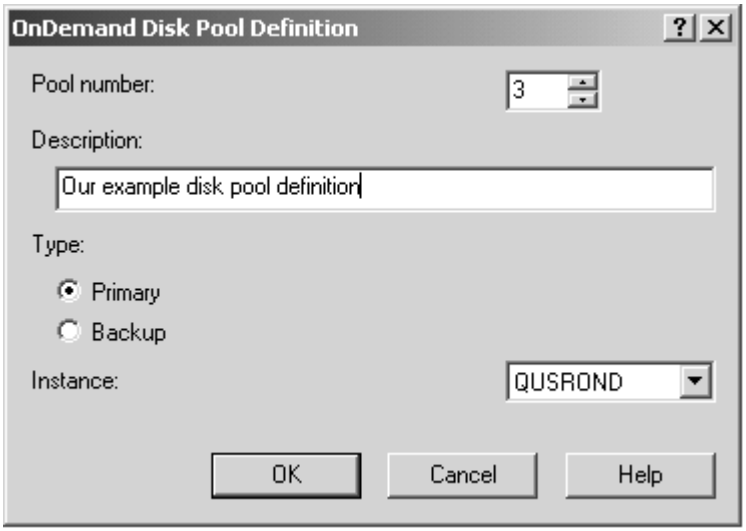


Figure 7. Disk Pool

Creating monitor definitions

1. To create a monitor definition, left click on the plus sign in front of Common Server Administration, right click on Monitor Definitions and select New Monitor.
2. When the New Monitor window opens up, the general page will be displayed. Here we need to specify the name and library of the output queue we would like to monitor. For this example we will specify the output queue as MONITORQ in library QRDARS. (However, your monitored output queues should not be placed in QRDARS library, since that library can be replaced during a software upgrade.) The job name is automatically filled in with the same name as the monitored output queue you specified. This can be changed, but for our example we will leave it as MONITORQ. You may also provide a brief description for the monitor definition.
3. When the monitor job selects a spooled file from the specified output queue for archiving, it needs to determine which application group to associate with the spooled file so that the file can be archived correctly. The only data available to the monitor are the attributes of the selected spooled file. The application group name must be derived from the contents of one of these attributes. Use the pull down to select the attribute that the monitor should use for comparing to find the application group with the matching name. For our example, we will use Spooled file name. You can also specify what the monitor should check next and last if it doesn't find a match on check first. In this example, we will leave check next and last blank.
4. You can check the box to have the monitor job determine the application name used or you can specify what the monitor should check for comparing to find the application name. In this example we will check the box and let the monitor job determine the application name.
5. You may also specify the job description, library and instance to be used for this monitor job. We will leave these values as the default for this example.
6. Once the general page has been completed you can click on the Start and End Methods tab. There are several options for starting and ending the monitor job. We will just accept the defaults for this example which are to start and end the

monitor manually. However you can specify whether the monitor will be started from a job scheduler or when a subsystem starts. You can also specify to have the monitor job end at a specific time of day, after a certain period of time, after all entries have been processed from the output queue, or you can specify the end time as a parameter when the monitor is started.

7. Next click on the Actions Taken on Spooled Files tab. This specifies what OnDemand should do with a spooled file once it has been successfully processed. The options are to have it deleted from the system or moved to an output queue you specify. In our example, we will have the spooled file moved to a processed output queue called PROCESSED in library QUSRRDARS. You must also specify where you want the spooled files to be sent should they fail to process. In our example, we will specify ERROR in library QUSRRDARS as our error queue.
8. Click on OK to create the monitor definition. See Figure 8, Figure 9 on page 62 and Figure 10 on page 63

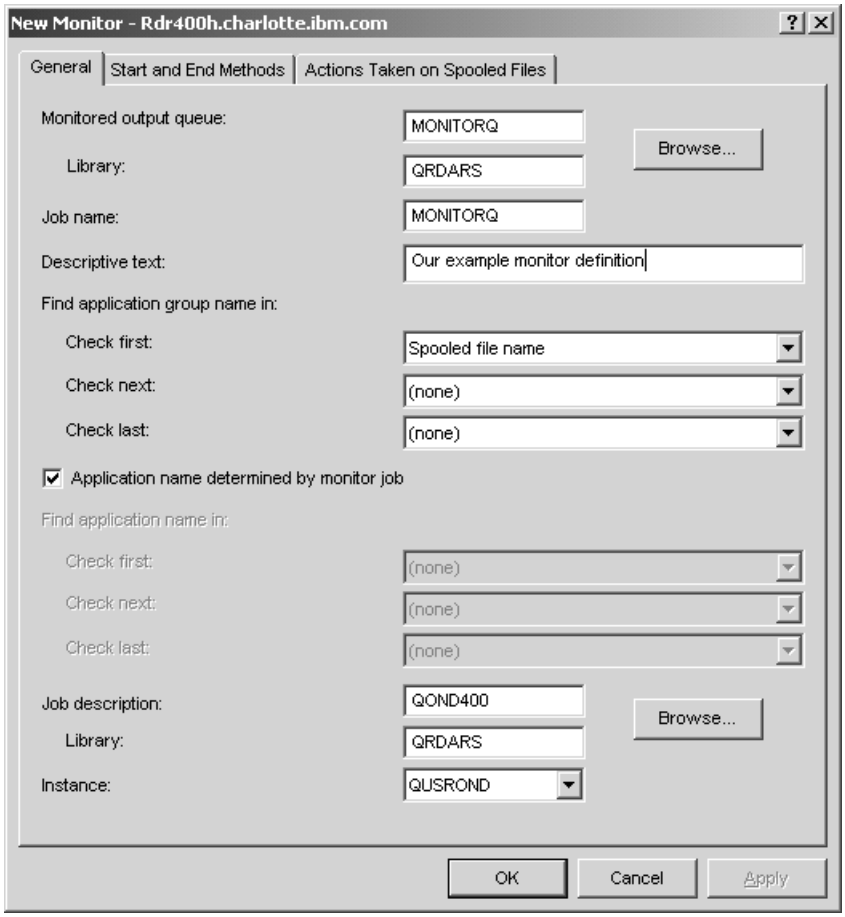


Figure 8. Monitor — General tab

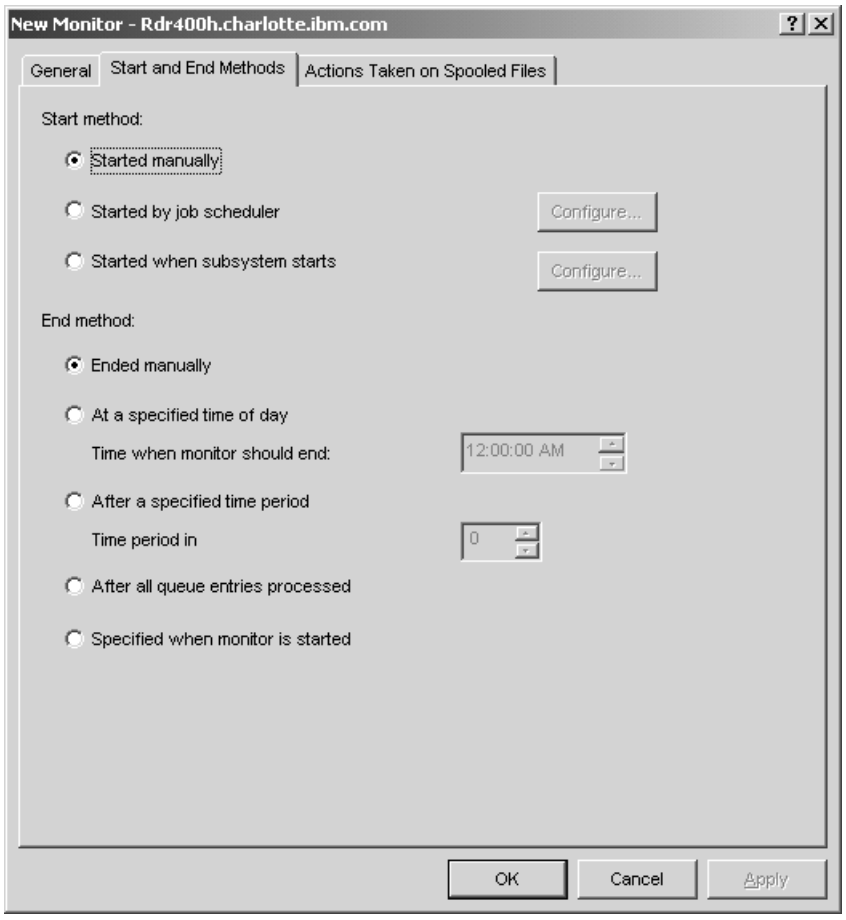


Figure 9. Monitor — Start and End Methods tab

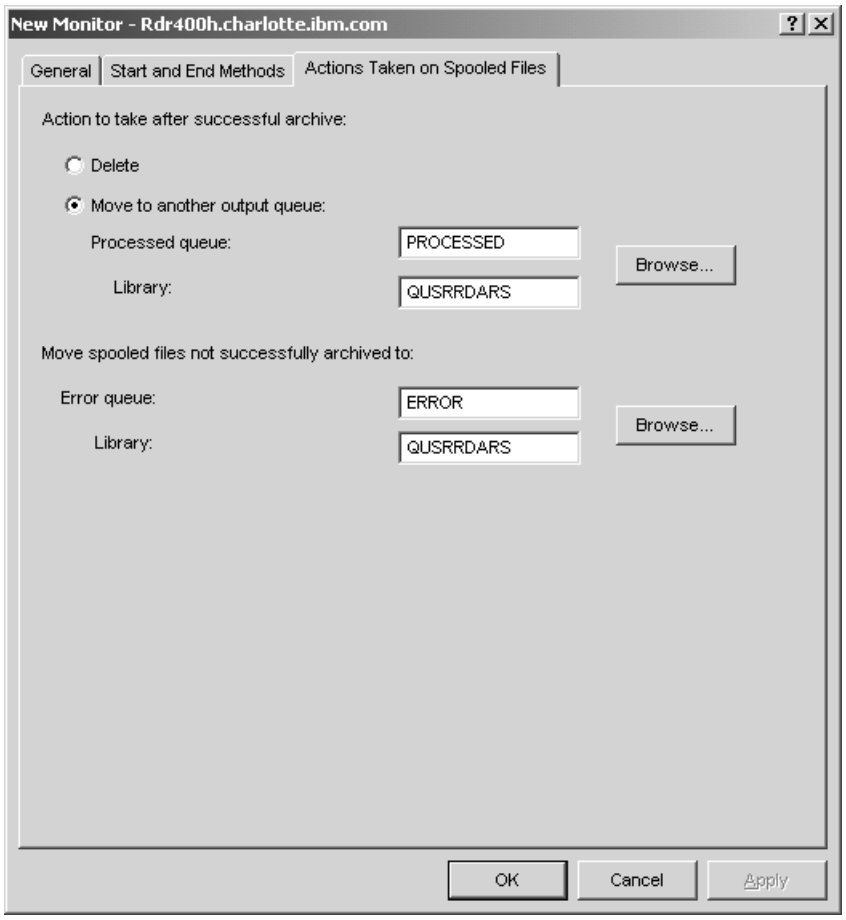


Figure 10. Monitor — Actions Taken on Spooled Files tab

Creating migration policies

1. To create a new migration policy, left click on the plus sign in front of Common Server Administration, right click on Migration Policies and select New policy.
2. First you must specify the name of your migration policy. The name can be up to 60 characters long, but must not duplicate the name of another policy within the same instance. For this example, we will choose POLICY1 as the new policy name. You may also provide a brief description of the migration policy you are creating.
3. You can choose whether or not you want to enable aggregation. If aggregation is enabled, you must select whether you want to close the aggregate only when the maximum size is reached or after a specified time period is reached. We will accept the default to enable aggregation and close the aggregate after the maximum size if reached. However, we will change the maximum size, by clicking on the up arrow, to 10,000 kilobytes.
4. You may also specify if a tape backup is requested. We will request a tape backup and select the Media type by clicking on the pull down and selecting 8MM.
5. The Instance is the logical server environment that will contain the archive data to which this migration policy will relate. For this example, we will leave the Instance as the default of QUSROND.

6. In our example, when a report is archived into OnDemand, a copy is stored on disk for 60 days. A copy is also maintained in optical storage for 5 years. This information will be used to set up our storage levels within our migration policy. For the first level, click on Add After, this opens the New Policy Level window. We must first assign a level identifier and this identifier must be unique within this policy. For this example we will use 0010 as our level identifier. You may also provide a brief description of this level. For the media, we will select Disk pool and the duration at this level should be 60 days. For the primary storage group, click on the pull down to see a list of available disk pool storage groups. We can select the one we created earlier called ASP01, then click OK. This will add our storage level in the storage levels window within our policy. To add the next level, select the level we just added and then click on Add After. Specify the level identifier as 0020 with a brief description. Our media is Optical and the number of days is 1825 (5 years). We will choose the optical storage group we created earlier, OPTSTG, as our primary storage group. Click OK, and you should now see both storage levels listed for our migration policy. In both our storage levels we did not select to Create backup copy. However, you can select this option and specify a backup storage group causing the migration process to create a duplicate copy of the archived data when it is moved to this level.
7. Click on OK to create the migration policy. Clicking OK will also create an OnDemand storage set of the exact same name. When you define an application group to OnDemand, you can select the storage set name that matches this migration policy, which will cause OnDemand to archive that application group data as defined here. See Figure 11.

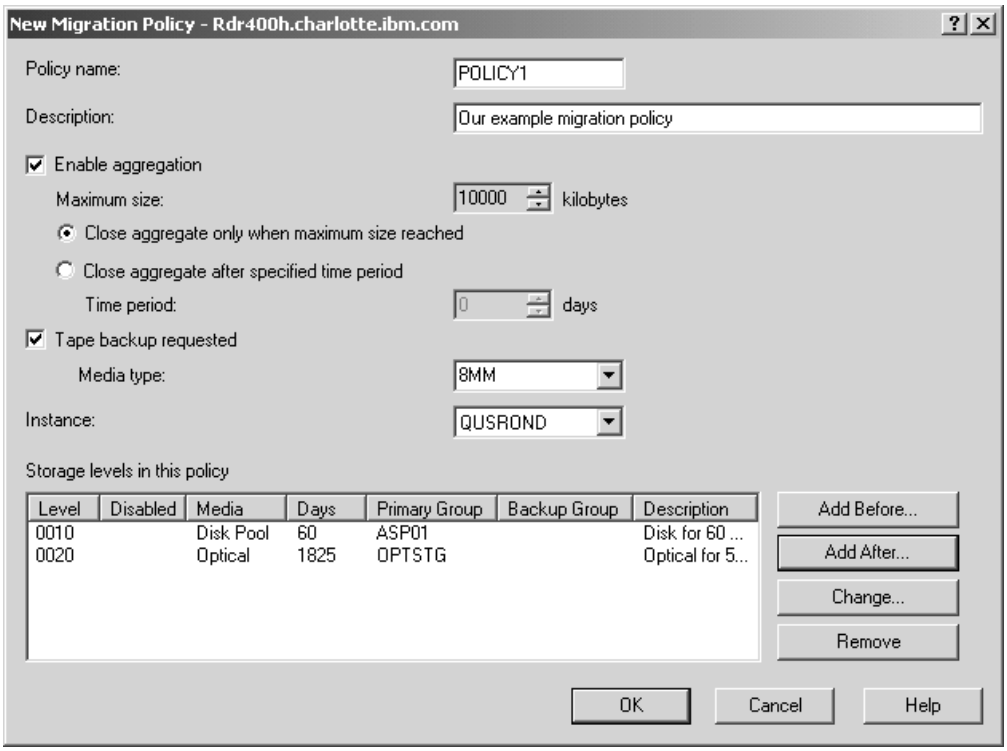


Figure 11. Migration policy

Adding users

The basics

In general, here is how you work in the administrative client to add a user:

1. In the OnDemand section of iSeries Navigator, right click Common Server Administration. In the resulting context menu, click Client Administrative Functions, then logon.
2. Choose a server and select Users.
3. Pick one of two ways to add a user.
4. Define the properties of the user by completing fields in the Add a User dialog box. Search for *user,adding* in the online help for more details.
5. Optionally add the user to groups. See the online help for more details.
6. Optionally add the user to application groups and set application group permissions. Search for *user,application group permission* in the online help for more details.
7. Optionally add the group to folders and set folder permissions. Search for *user,folder permission* in the online help for more details.
8. When finished, add the user by clicking OK in the Add a User dialog box.

Note: To add a user, the logon userid must be a user with create users authority, a user administrator, or a system administrator.

Users who need to run server commands such as those listed in Appendix A, "Command reference" on page 133 or server APIs from QSHELL such as those listed in Appendix D, "API and user exit reference" on page 145 need to have QRDARSADM as the group profile (or a supplemental group) in their OS/400 user profile.

Choose a server

1. On the left side of the administrator window, click the name of the server to which you want to add the user.
2. Expand the areas of the server. Double click the server name or click the + (plus) to the left of the server name.
3. Select Users.

Two ways to add a user

This section explains how to add a user by using the New User command. You can also add a user by copying an existing user definition.

New User command: From the File menu, select New User to open the Add a User dialog box.

Copy command: You can use the copy command to add a user. In the User ID list, point to the user that you want to copy and click the right mouse button. Select Copy from the pop-up menu to open the Add a User dialog box. The fields in the dialog box contain information copied from the user you selected. At a minimum, you need to change the User ID (userids must be unique to the server).

Adding the user

In the Add a User dialog box, click OK. The administrative client adds the user to the database and returns to the main window.

Examples

First, let's review the requirements of the users that need to access the telephone bill reports.

- We need to define a set of userids for the customer service department. We'll identify one user as a user administrator. The user administrator can add other users.
- We need to define a userid for Customer XYZ.

On our system, users obtain permissions from groups. That means we don't add users to application groups and folders. When we define our groups, we'll add users to the groups.

Note: We're not going to document how to add all of the users to our system. We'll add two users. You can repeat the steps to add the others.

Adding the customer service users

Let's use the New User command to add a user.

1. First, point to Users and click the right mouse button. From the pop-up menu, select New User to open the Add a User Dialog box.
2. In the User ID field, enter the name of the user: *CSR1* If your OnDemand userids are linked to your OS/400 user profiles, then this userid must also exist as an OS/400 user profile.
3. Accept the UID generated by OnDemand.
4. Set the user's initial password to the userid. In the Password field, enter: *CSR1* If your OnDemand userids are linked to your OS/400 user profiles, then this password will be ignored.
5. In Verify Password, enter: *CSR1*
6. In the Description field, enter: *Customer Service Representative*
7. Under User Type, select User Administrator. The user will be able to maintain userids on the system.
8. Click the User Information tab.
9. Complete the fields on the User Information page, such as the Name, Department, and Phone Number.
10. Click the General tab.

    The user will obtain application group and folder permissions from a group. We'll add the user to the group when we define the group. Therefore, we do not need to add the user to application groups, folders, or groups at this time.

    Figure 12 on page 67 shows the Add a User dialog box with the basic information completed.

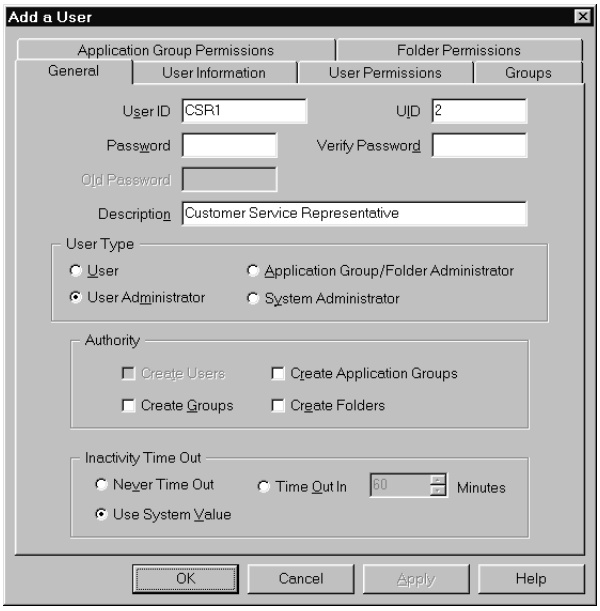


Figure 12. Add a User dialog box - basic information

11. At this point, the properties of the user meet the requirements. Click OK to add the user.

Let's use the Copy command to add another user.

1. Point to CSR1 and click the right mouse button. From the pop-up menu, select Copy to open the Add a User Dialog box.
2. In the User ID field, replace *CSR1* with *CSR2* If your OnDemand userids are linked to your OS/400 user profiles, then this userid must also exist as an OS/400 user profile.
3. Accept the UID generated by OnDemand.
4. Set the user's initial password to the userid. In the Password field, enter: *CSR2* If your OnDemand userids are linked to your OS/400 user profiles, then this password will be ignored.
5. In the Verify Password, enter: *CSR2*
6. Under User Type, select User.
7. Click the User Information tab.
8. Replace the information in the fields on the User Information page.
9. Click the General tab.

   Figure 13 on page 68 shows the Add a User dialog box with the basic information completed.

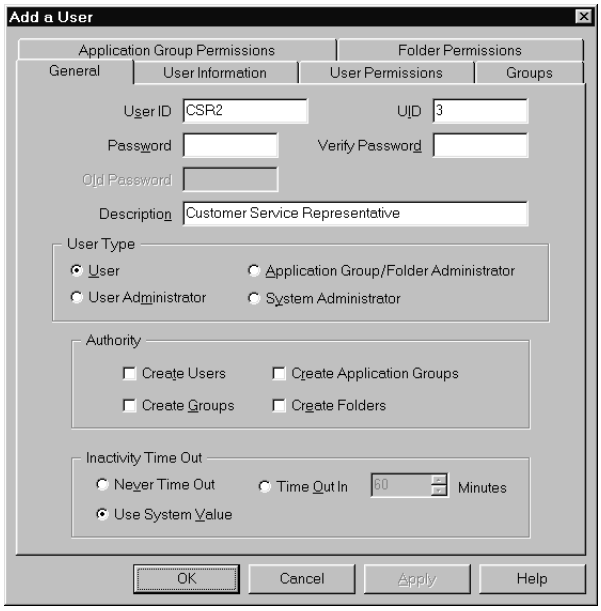


Figure 13. Add a User dialog box - basic information

10. At this point, the properties of the user meet the requirements. Click OK to add the user.

Adding groups

In general, here is how you work in the administrative client to add a group:

1. Choose a server and select Groups.
2. Pick one of two ways to add a group.
3. Define the properties of the group by completing fields in the Add a Group dialog box. Search for *group,adding* in the online help for more details.
4. Optionally assign a group owner.
5. Optionally add users to the group.
6. Optionally add the group to application groups and set application group permissions. Search for *group,application group permissions* in the online help for more details.
7. Optionally add the group to folders and set folder permissions. Search for *group,folder permissions* in the online help for more details.
8. When finished, add the group by clicking OK in the Add a Group dialog box.

Note: To add a group, the logon userid must be a user with create groups authority or a system administrator.

Choose a server

1. On the left side of the window, click the name of the server to which you want to add the group.
2. Expand the areas of the server. Double click the server name or click the + (plus) to the left of the server name.
3. Select Groups.

Two ways to add a group

This section explains how to add a group by using the New Group command. You can also add a group by copying an existing group definition.

New Group command

From the File menu, select New Group to open the Add a Group dialog box.

Copy command

You can use the copy command to add a group. In the Name list, point to the group that you want to copy and click the right mouse button. Select Copy from the pop-up menu to open the Add a Group dialog box. The fields in the dialog box contain information copied from the group you selected. At a minimum, you need to change the group name (group names must be unique to the server).

Note: These are not the same as group profiles in OS/400. However, the names may match if you find that easier to maintain.

Adding users

You can add one or more users to the group. Complete the following steps to add a user to a group.

1. From the List of Users list, select the user.
2. Click Add. The administrative client moves the user to the Users in the Group list.

   To remove a user from the group, select the user in the Users in the Group list and click Remove.

Adding the group

In the Add a Group dialog box, click OK. The administrative client adds the group to the database and returns to the main window.

Examples

First, let's review the requirements of the groups that need to access the telephone bill reports.

- Customer service group. Users that belong to the group can open the telephone bill report folder and query documents stored in the telephone bill report application group. When we define the report to the system, we'll add the group to the application group and the folder. We need to identify a group owner. The group owner can add new customer service users to the group and remove users from the group.
- Customer XYZ group. Users that belong to the group can also open the telephone bill report folder and query documents stored in the telephone bill report application group. However, we'll limit access to documents that contain a specific customer name and account number. When we define the report to the system, we'll add the group to the application group and the folder and set up the necessary restrictions.
- Users. We add the users that we defined in "Adding users" on page 65 to the groups.

Adding the customer service group

Let's use the New Group command to add the group.

1. First, point to Groups and click the right mouse button. From the pop-up menu, select New Group to open the Add a Group Dialog box.
2. In the Name field, enter the name of the group: *CustomerService*

3. Accept the GID generated by OnDemand.
4. From the Group Owner list, select the user that can add users to and remove users from the group: CSR1

   Note: The user that you select does not obtain permissions from the group unless you add the user to the group (see step 6). However, a group owner can add their userid to the group at any time.
5. In the Description field, enter: *Access to Telephone Bill Reports; call Leonard Little, x90565, for more information*
6. From the List of Users list, select and add users to the group: CSR1, CSR2, CSR3, CSR4, and CSR5

   Figure 14 shows the Add a Group dialog box with the basic information completed.

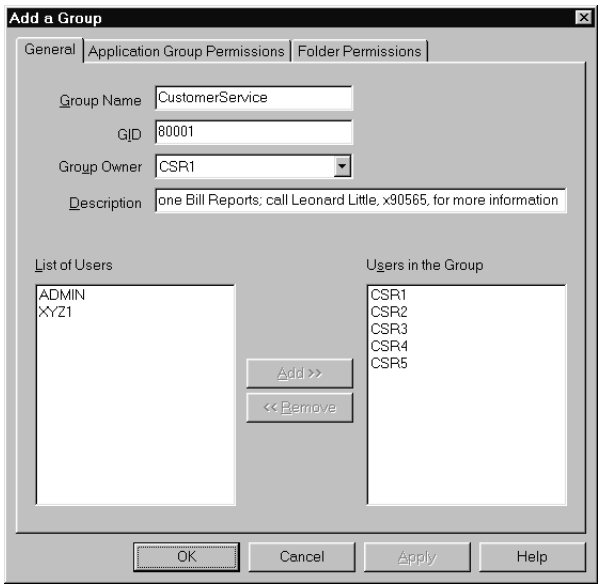


Figure 14. Add a Group dialog box

7. At this point, the properties of the group meet the requirements. Click OK to add the group.

Adding the Customer XYZ group

Let's use the Copy command to add the group.

1. Point to Customer Service and click the right mouse button. From the pop-up menu, select Copy to open the Add a Group Dialog box.
2. In the Name field, replace *CustomerService* with *CustomerXYZ*
3. Accept the GID generated by OnDemand.
4. Since we want a system administrator to maintain the group, do not assign a group owner. In the Group Owner list, replace CSR1 with *NONE.
5. Replace the contents of the Description field with: *Access to Telephone Bill Reports by Customer XYZ*
6. From the Users in the Group list, remove CSR1, CSR2, CSR3, CSR4, and CSR5.
7. From the List of Users list, add XYZ1.

   Figure 15 on page 71 shows the Add a Group dialog box with the basic information completed.

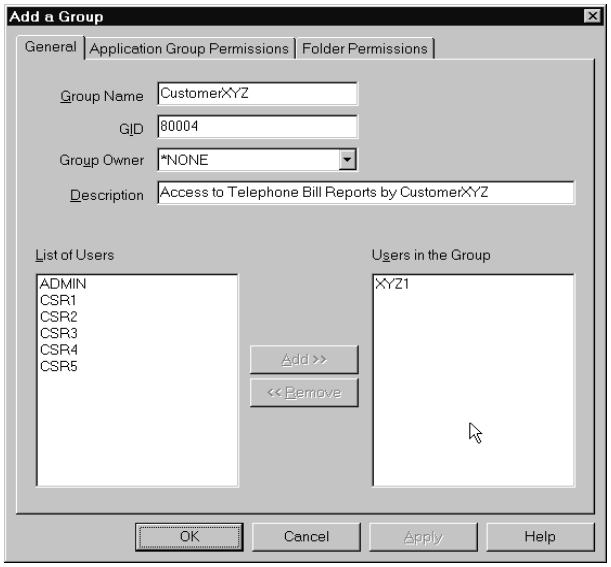


Figure 15. Add a Group dialog box

8. Click OK to add the group.

Adding server printers

In general, here is how you work in the administrative client to add a server printer:

1. Choose a server and select Printers.
2. Pick one of two ways to add a server printer.
3. Define the properties of the server printer by completing fields in the Add a Printer dialog box. Search for *printer,adding* in the online help for more details.
4. When finished, add the server printer by clicking OK in the Add a Printer dialog box.

Notes:

1. To add a server printer, the logon userid must be a system administrator.
2. Server print is supported for spooled file data types of SCS, SCS-extended, AFPDS, and LINE.

Choose a server

1. On the left side of the main window, click the name of the server to which you want to add the server printer.
2. Expand the areas of the server. Double click the server name or click the + (plus) to the left of the server name.
3. Select Printers.

Two ways to add a server printer

This section explains how to add a server printer by using the New Printer command. You can also add a server printer by copying an existing server printer definition.

New Printer command

From the File menu, select New Printer to open the Add a Printer dialog box. Figure 16 on page 72 shows an example of the Add a Printer dialog box.

Copy command

You can use the copy command to add a server printer. In the Name list, point to the server printer that you want to copy and click the right mouse button. Select Copy from the pop-up menu to open the Add a Printer dialog box. The fields in the dialog box contain information copied from the server printer you selected. At a minimum, you need to change the name (printer names must be unique to the server).

Adding the server printer

In the Add a Printer dialog box, click OK. The administrative client adds the server printer to the database and returns to the main window.

Examples

We're going to add a server printer that can be used as the default server printer for the application that we'll define in "Adding a report" on page 73. The physical printer resides in the customer service department. By default, when users of the telephone bill report application select a document and choose the server printer command, OnDemand sends the document to this print device.

Adding the server printer

Let's use the New Printer command to add the server printer.

1. First, point to Printers and click the right mouse button. From the pop-up menu, select New Printer to open the Add a Printer dialog box.
2. In the Name field, enter the name of the server printer: *Customer Service printer*
3. In the Description field, enter: *Customer Service printer; default server printer for the telephone bill report application*
4. In the Server Queue Name, enter: *ip60cs* If no library name is entered (as in the example), OnDemand assumes *LIBL; otherwise, enter a value such as QUSRSYS/PRT04 to identify both the library name and the output queue name.
5. Accept the default Printer Type of *Printer*

   Figure 16 shows the completed Add a Printer dialog box.

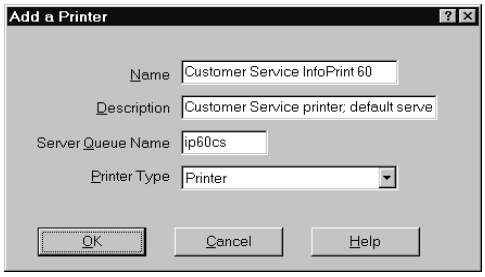


Figure 16. Add a Printer dialog box

6. At this point, the properties of the server printer meet the requirements. Click OK to add the server printer.

Adding storage sets

A storage set is automatically added each time you create a migration policy. The name of the storage set that is created is an exact match of the name of the migration policy. See "Creating migration policies" on page 63 for details on migration policies.

Adding a report

In general, here is how you work in the administrative client to define a report to the system:

1. Choose a server. On the left side of the main window, click the name of the server on which you want to define the report.
2. Select the area. First expand the areas of the server. Double click the server name or click the + (plus) to the left of the server name. Then select the area.

   When you define a report to the system, you typically add an application group, an application, and a folder.
   - The application group identifies database and storage management information.
   - The application identifies viewing, indexing, loading, and printing information.
   - The folder provides users the ability to search for, retrieve, view, and print report data.
3. Pick one of two ways to add an application group, an application, and a folder. For example, use one of the following methods to add an application group:
   - Add a new application group. From the File menu, select New Application Group to open the Add an Application Group dialog box.
   - Copy an existing application group. In the Name list, point to the application group that you want to copy and click the right mouse button. Select Copy from the pop-up menu to open the Add an Application Group dialog box. The fields in the dialog box contain information copied from the application group you selected. At a minimum, you need to change the Name (application group names must be unique to the server).
4. Define the properties of the application group, application, and folder.
5. When finished, add the application group, application, and folder by clicking OK in the add dialog box.

Note: To add an application group, the logon userid must be a system administrator, an application group/folder administrator, or a user with create application groups authority.

To add an application, the logon userid must be a system administrator, an application group/folder administrator, or a user with create application groups authority.

To add a folder, the logon userid must be a system administrator, an application group/folder administrator, or a user with create folders authority.

An example

This section describes how to define a sample telephone bill report to the system.

About the report

A telephone bill report typically contains hundreds of pages of line data. The report is logically segmented into statements. Users search for statements using a date and any combination of account number and customer name.

Most queries about a statement occur in the first 60 days after it is mailed to the customer. Little or no activity occurs a year after a statement is generated. For legal

reasons, we need to keep a statement for five years. We expect the system to maintain index information in the most efficient way possible.

Two groups of users need to access the telephone bill reports. The customer service department is responsible for handling queries from customers. They answer questions about the statements, attach annotations to statements, and reprint and fax a replica of original statements. As part of our enhanced customer service initiative, we're allowing Customer XYZ to access the system and retrieve and view their statements.

About the application group

Before we add the application group, let's review the database and storage management requirements.

Database requirements: Database requirements can be grouped in two categories: database tables and the database fields.

Database tables:

- A database table contains index data from one or more reports
- An annotation field is not required in the database
- The date field is the segment field for the application group

Database fields:

- Three database fields: account number, customer name, and report date
- Account number is the index, data type string, 16 bytes. However, we plan to store only 13 bytes in the database, removing the - (dash) characters from the account number string.
- Customer name is a filter, data type string, 30 bytes
- Report date is a filter, data type date

Storage management requirements: The storage management requirements determine where, how, and how long OnDemand maintains the report and index data.

- Maintain a report for five years
- Copy documents to disk and maintain them for 60 days
- Copy documents to archive storage when the report is loaded into the system
- Delete a table of index data at a time

Adding the application group

Let's use the New Application Group command to add the application group.

1. First, point to Application Groups and click the right mouse button. From the pop-up menu, select New Application Group. The administrative client opens the Add an Application Group dialog box.

   The pages of the Add an Application Group dialog box organize information about the application group into sections. The tabs show which page you are on: General, Message Logging, Storage Management, Permissions, Field Definition, and Field Information.
2. We'll start by completing the General page. In the Name field, enter the name of the application group.
3. In the Description field, enter up to 120 characters of descriptive information about the application group.

4. Based on our database requirements, we can accept the recommended defaults for the rest of the fields on the General page. (You can click Advanced to see the default options.) Search for *application group,general* in the online help for more details.

   Figure 17 shows the completed General page.

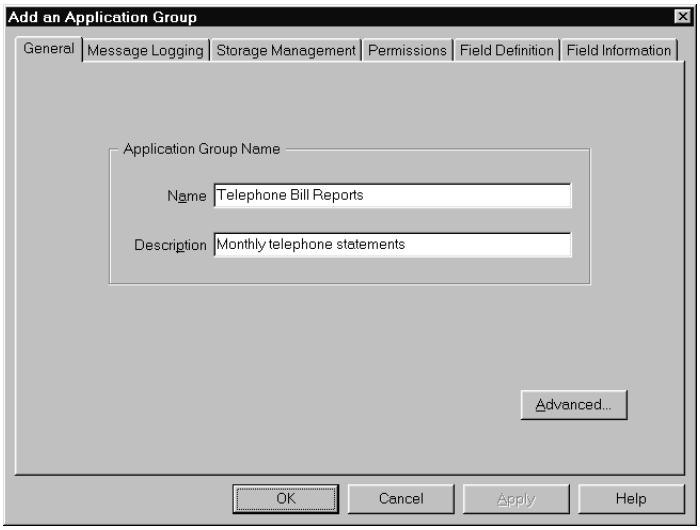


Figure 17. General page

5. Select the Message Logging tab to specify the types of application group messages that we want OnDemand to save in the system log. Search for *application group,message logging* in the online help for details about the options that you can specify.
6. Select Retrieval, Database Queries, and Server Printing. Clear all of the other check boxes. We'll use the information in the system log messages to generate system usage reports.

   Figure 18 shows the completed Message Logging page.

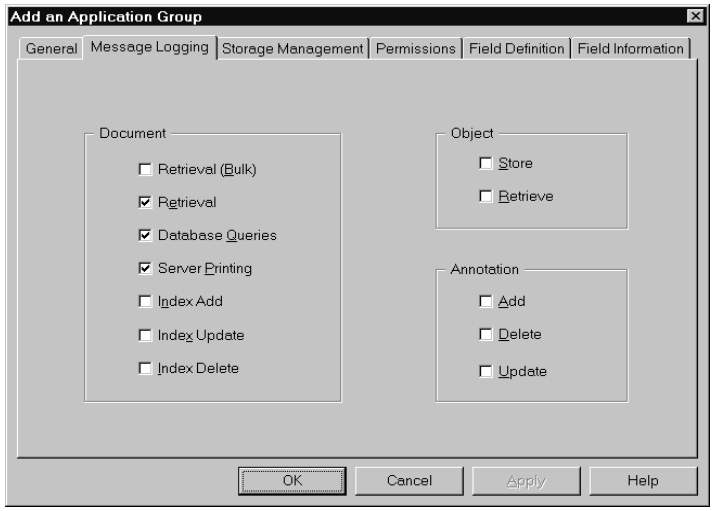


Figure 18. Message Logging page

7. Select the Storage Management tab to provide information that OnDemand uses to manage data stored in the application group. Search for *application group,storage management page* in the online help for details about the options that you can specify.
8. From the Storage Set Name list, select the name of the storage set that OnDemand will use to maintain the report on disk storage and in archive storage. We defined the storage set and migration policy in "Adding storage sets" on page 72.
9. Under Cache Data, select `Cache Data for __ Days` and enter `60` in the space provided. We want OnDemand to maintain the report on disk storage for 60 days. After that time, when a user views or prints the report, OnDemand retrieves the report from archive storage.
10. Under Life of Data and Indexes, select `Expire in ____ Days` and enter `1825` in the space provided. This is the number of days (1825, or five years) that we want OnDemand to maintain index data, documents, and resources related to the report. After that number of days, OnDemand can delete the report from the system.
11. Because we want OnDemand to delete a table of index data at a time, from the `Expiration Type` list, select `Segment`.

    Figure 19 shows the completed Storage Management page. Although our example shows the selection of an `Expiration Type of Segment`, the recommended selection for this field is `Load`. (You can click Advanced to see other, default options.)

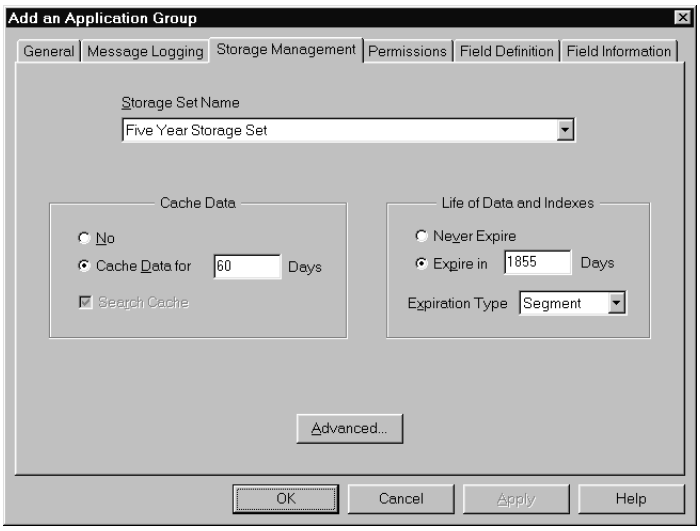


Figure 19. Storage Management page

12. Select the Permissions tab to specify the types of report and application group functions that users can perform. For example, we can let users query report data, create logical views, print and fax pages of the report, and maintain the application group. We can specify default permissions and permissions for specific groups and users. By default, the person that creates the application group is given all application group permissions; no other users can access report data or maintain the application group. On our system, other users obtain permissions from a group. We added groups in "Adding the customer service group" on page 69 and "Adding the Customer XYZ group" on page 70. Search for *application group,permissions page* in the online help for details about the options that you can specify.

13. First, add the customer service group. From the User/Groups list, select +CustomerService.
14. Select the Access check box.
15. Click Add.
16. Next, add the Customer XYZ group. From the User/Groups list, select +CustomerXYZ.
17. In the Annotation area, clear the View check box. OnDemand also clears the Add check box. Customer XYZ users can view, print, fax, and copy documents but do not have permission to use the annotation tools provided by OnDemand.
18. Click Add.

    Figure 20 shows the completed Permissions page.

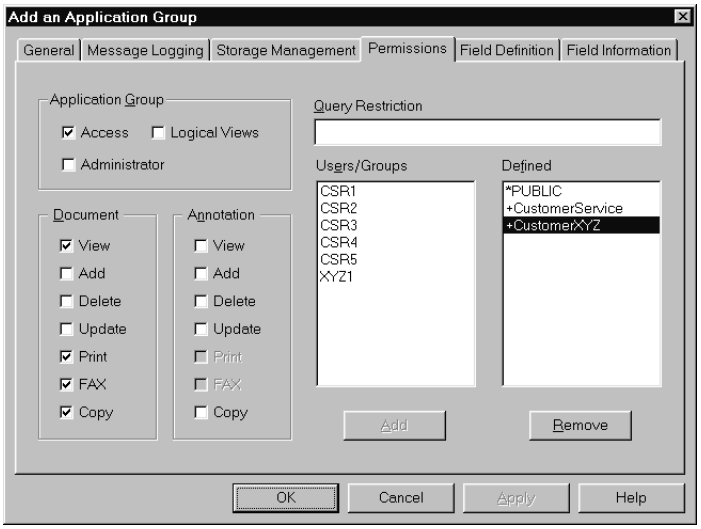


Figure 20. Permissions page

19. Select the Field Definition tab to define the database fields. We need to define three database fields for the telephone bill report:
    - Account number
    - Customer name
    - Report date

    To define a database field, enter the name of the field in the Database Field Name field and click Add. Search for *application group,field definition* in the online help for more details.

    Figure 21 on page 78 shows the completed Field Definition page.

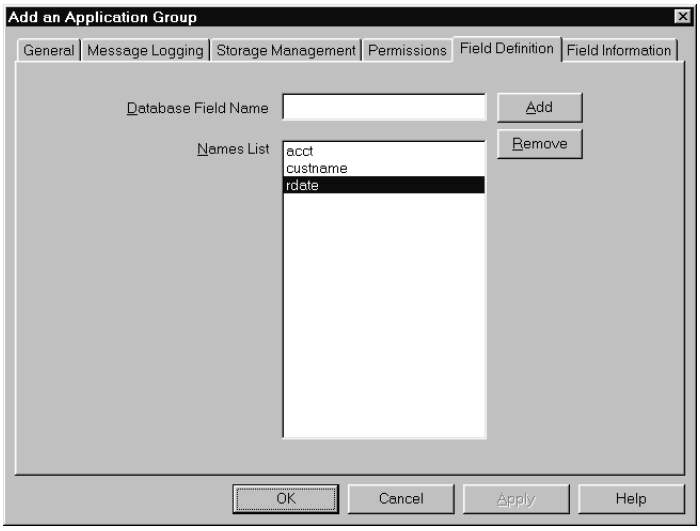


Figure 21. Field Definition page

20. Select the Field Information tab to define the attributes of the database fields. Search for *application group,field information page* in the online help for details about the options that you can specify.
21. First, define field information for the account number field.
    - From the Type list, select Index
    - Enter the string length (13)

    Figure 22 shows the Field Information page for the account number field.

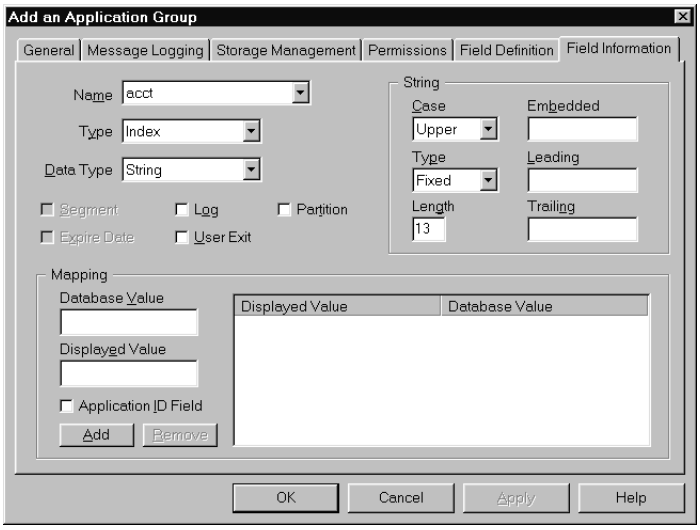


Figure 22. Field Information page

It is highly recommended that you ALWAYS define an Application ID Field on the Field Information page. (Figure 22 shows where you would define one.) You may not think you need one at the time you create your application group, but for flexibility in the future, you may need one later. The important thing to note is that an application ID field CANNOT be added to the application group AFTER it is created; it must be part of the original application group definition. Also note that the application ID field can be

hidden so that end-users are not confused by it if it does not provide value to them. Refer to online help for more information about the Application ID Field.

22. Next, define field information for the customer name field.
    - From the Name list, select custname
    - From the Type list, select Variable
    - In the Length field, enter 30

    Figure 23 shows the Field Information page for the customer name field.

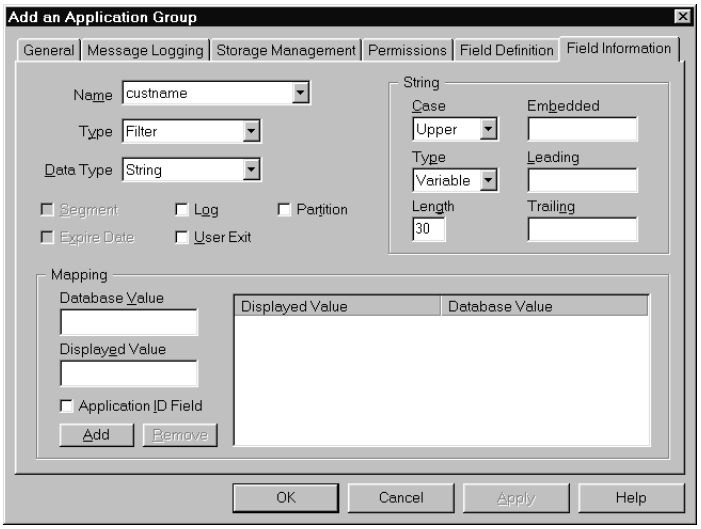


Figure 23. Field Information page

23. Next, define field information for the report date field.
    - From the Name list, select rdate
    - From the Data Type list, select Date
    - Select the Segment check box

    Figure 24 on page 80 shows the Field Information page for the report date field.

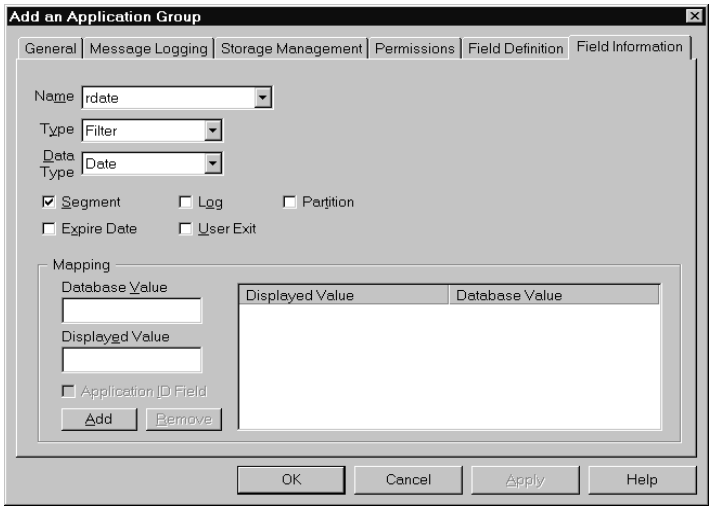


Figure 24. Field Information page

24. At this point, the properties of the application group meet the requirements. Click OK in the Add an Application Group window. The administrative client adds the application group to the database and returns to the main window.

About the application

Before we add the application, let's review the viewing, indexing, and loading requirements.

Viewing requirements:

- Source data stored in OnDemand as AFP data
- Format data into pages and enhance the appearance with images and fonts
- Retrieve statements of one or more pages
- Define a default printer for the application. The default printer is where OnDemand sends documents when users select the server print command.

Indexing requirements:

- Source data is EBCDIC, code page 500
- Segment report into groups of pages, one statement in each group
- Identify the beginning of a statement using:

  Skip-to-channel one (X'F1')

  PAGE 1 (X'D7C1C7C5404040F1')
- Generate three indexes for each statement: statement date, account number, and customer name
- Collect resources

Loading requirements:

- Compress and store data in the most efficient method possible
- Application group database field names and index names match
- Date format is Mth d, yyyy
- Remove embedded - (dash) character from account number before storing value in the database

Adding the application

Let's use the New Application command to add the application.

1. First, point to Applications and click the right mouse button. From the pop-up menu, select New Application to open the Add an Application dialog box.

   The pages of the Add an Application dialog box organize information about the application into sections. The tabs show which page you are on: General, View Information, Indexer Information, Load Information, Logical View Fields, Logical Views, and Print Options.
2. We'll start by completing the General page. In the Name field, enter the name of the application.
3. In the Description field, enter information about the application.
4. Click Select to open the Application Groups dialog box.
5. From the Names list, select Telephone Bill Reports.
6. Click OK.

   Figure 25 shows the completed General page. Search for *application,general page* in the online help for details about the options that you can specify.

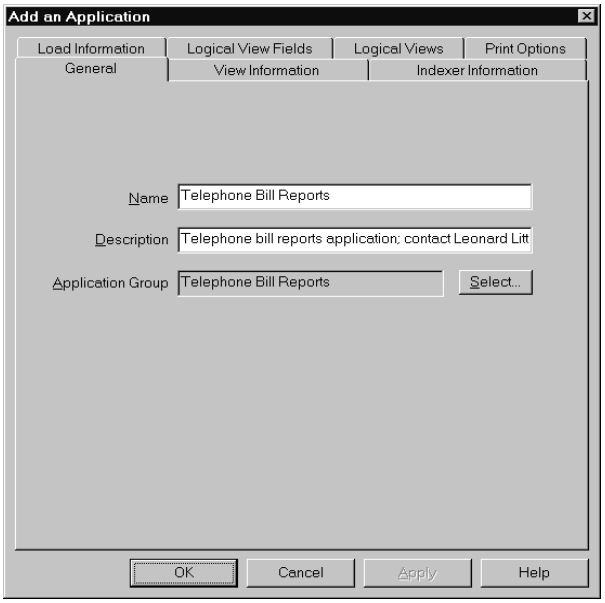


Figure 25. General page

7. Select the View Information tab to specify information needed by OnDemand client programs to display the telephone bill report. This information is also used by the indexing program. Search for *application,view information page* in the online help for details about the options that you can specify.
8. From the Data Type list, select Line. Note that although this example shows LINE data type, most OS/400 spooled files will be either SCS or AFP. Also note that AFP overlays are supported only when using the AFP data type. If you have SCS spooled files with AFP overlays and you need to keep the overlay with the print data for viewing, printing, and so forth, then your printer file needs to be changed to *AFPDS prior to creating the spooled file. Then, when defining the report to OnDemand, the application's Data Type field would be set to AFP.
9. In the RECFM area, select Fixed. The report contains fixed length records, 133 bytes in length.

   Figure 26 on page 82 shows the View Information page.

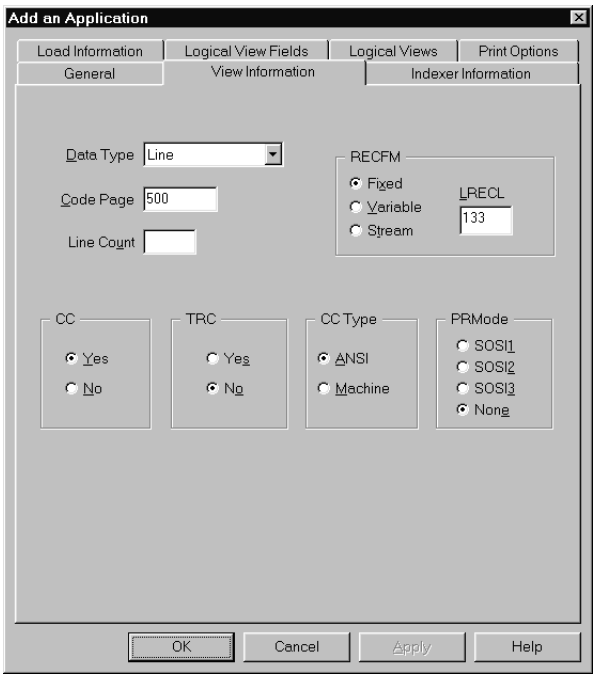


Figure 26. View Information page

10. Select the Indexer Information tab. Search for *application,indexer information* in the online help for details about the options that you can specify.
11. From the Indexer list, select OS/400 indexer.
12. We'll process a sample report using the graphical indexer. In the Parameters Source area, select Sample Data.
13. Click Modify to open the Open dialog box to select a file that contains a sample of the actual report data.
14. Select OS/400 Spooled File to work with spooled files on the server. Identify how to find your spooled file by entering a user profile or output queue name, and then click Retrieve List to generate a list of spooled files to choose from. Once you select a specific spooled file, OnDemand copies that sample data to your workstation for you to use for indexing.

    Figure 27 on page 83 shows the report window.

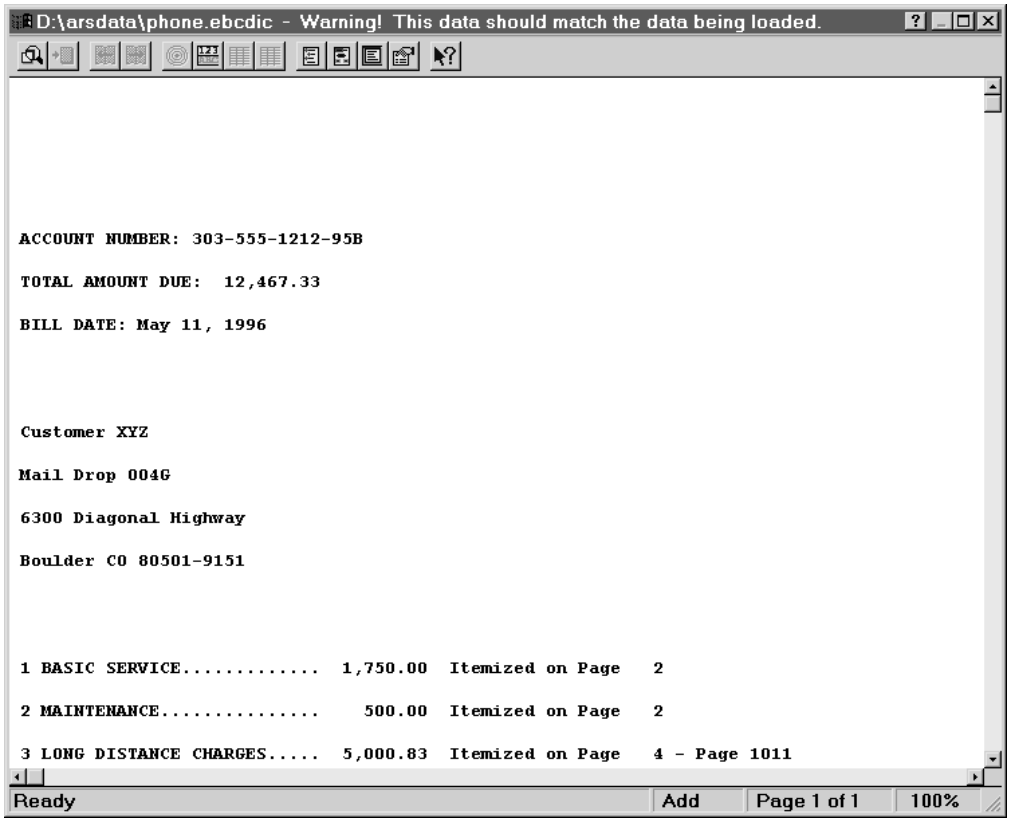


Figure 27. Report window

15. Define trigger number one. First, select any blank column in the first record. Click the Add a Trigger icon to open the Add a Trigger dialog box. (See online help for a description of trigger fields.)
16. In the Columns to Search area, select Carriage Control.

    Figure 28 shows the completed Add a Trigger dialog box for Trigger1.

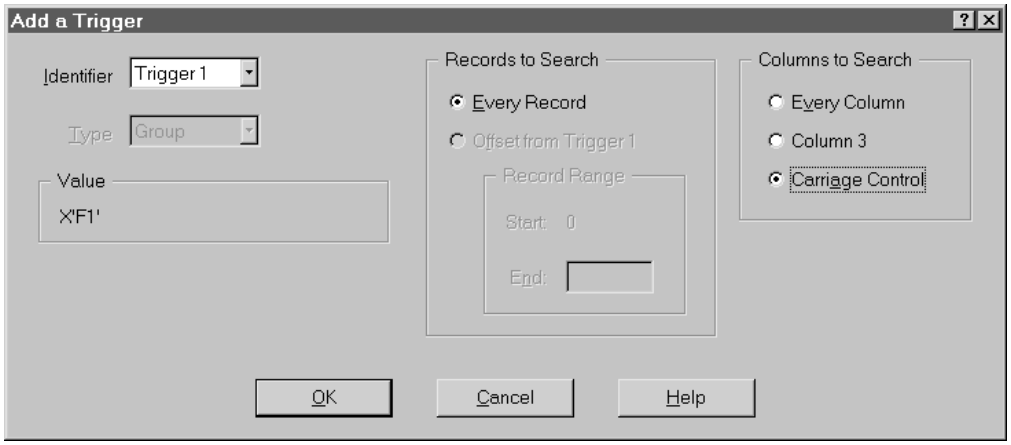


Figure 28. Add a Trigger dialog box

17. Click OK to add the trigger.
18. Define trigger number two. First, select the string PAGE 1. Click the right mouse button. From the pop-up menu, select Trigger to open the Add a Trigger dialog box.

19. Click OK to add the trigger.
20. Define field number one. First, select the string Customer XYZ and enough blank characters to the right of the string to reserve enough space to hold the largest index value (30 characters) the field can contain. (The selected string length guide, which appears above the field, displays the number of selected characters.) Then click the right mouse button. From the pop-up menu, select Field to open the Add a Field dialog box.
21. From the Trigger list, select Trigger2.

    Figure 29 shows the completed Add a Field dialog box for Field1.

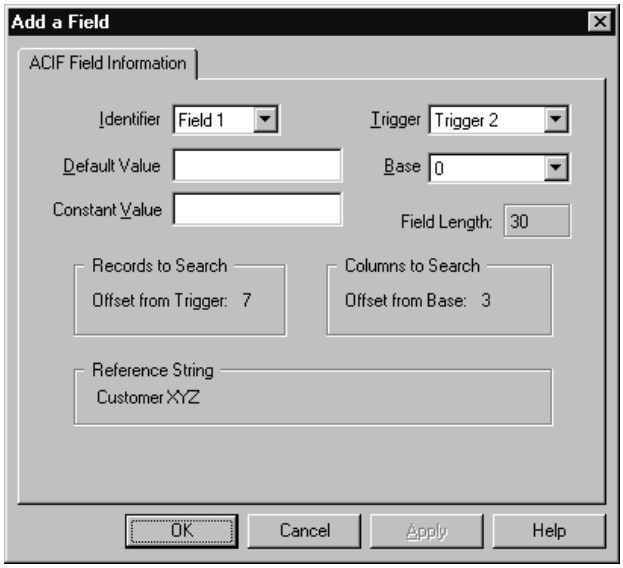


Figure 29. Add a Field dialog box

22. Click OK to add the field.
23. Define field number two. First, select the string May 11, 1996. Then click the right mouse button. From the pop-up menu, select Field to open the Add a Field dialog box.
24. From the Trigger list, select Trigger2.
25. Click OK to add the field.
26. Define field number three. First, select the string 303-555-1212-95B. Then click the right mouse button. From the pop-up menu, select Field to open the Add a Field dialog box.
27. From the Trigger list, select Trigger2.
28. Click OK to add the field.
29. Define the indexes. First, clear any selected areas of the report. Click the Add an Index icon to open the Add an Index dialog box.
30. From the Attribute list, select custname.
31. In the Break area, select No.
32. In the Fields list, double-click Field1.

    Figure 30 on page 85 shows the Add an Index dialog box for Index1.

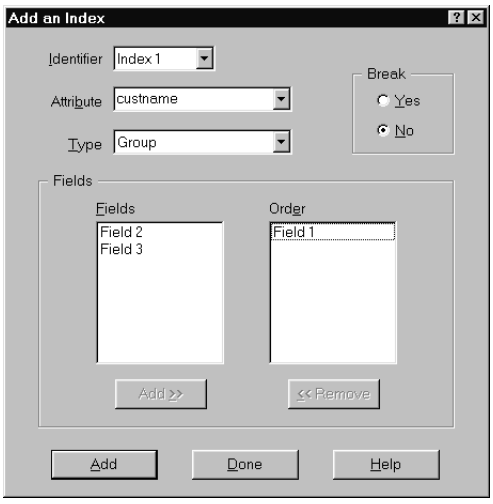


Figure 30. Add an Index dialog box

33. From the Attribute list, select rdate
34. In the Break area, select No.
35. In the Fields list, double-click Field2.
36. From the Attribute list, select acct.
37. In the Fields list, double-click Field3.
38. Click Done to close the Add an Index dialog box.
39. Close the report window, saving the changes.
40. Select the Load Information tab to specify information that OnDemand uses to process the index data before storing it in the database. Search for *application,load information page* in the online help for details about the options that you can specify.
41. In the Application Group DB Name list, select rdate. We need to specify the format of the date as it appears in the report.
42. From the format list, select %b %e, %Y

    Figure 31 on page 86 shows the Load Information for the report date field.

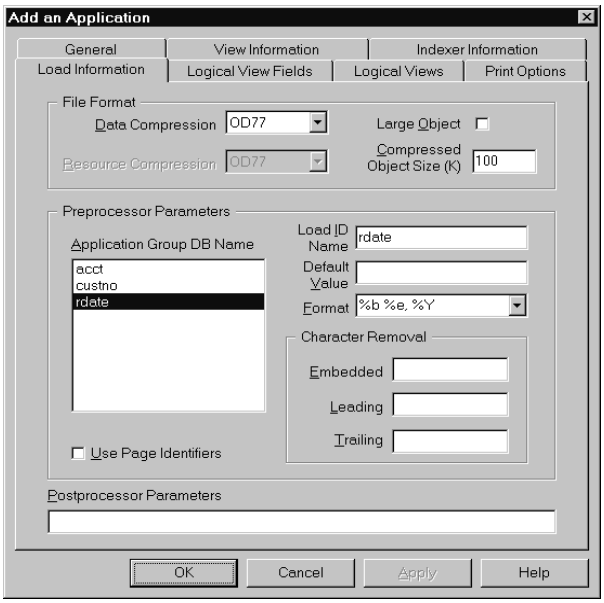


Figure 31. Load Information page

43. In the Application Group DB Name list, select acct. To conserve space in the database, we want OnDemand to remove the - (dash) character from index values before storing the values in the database. Also, it is best to define your numeric index fields (including date fields) so that leading, trailing, and embedded blanks, dashes, currency symbols, thousands separators, etc. are removed by OnDemand prior to being stored. Edits on these fields are sometimes strict and can cause a load to fail if non-numeric characters are found within the fields being defined.
44. In the Embedded field, enter the - (dash) character.

    Figure 32 shows the Load Information for the account number field.

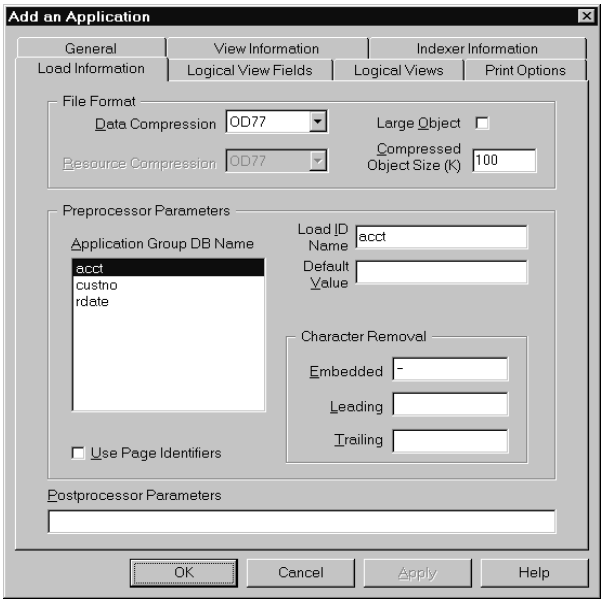


Figure 32. Load Information page

45. Select the Print Options tab to provide information that OnDemand uses to print the report. Search for *application,print options page* in the online help for details about the options that you can specify.
46. From the Default Server Printer list, select Customer Service printer. This is the printer we defined in "Adding server printers" on page 71.

    Figure 33 shows the completed Print Options page.

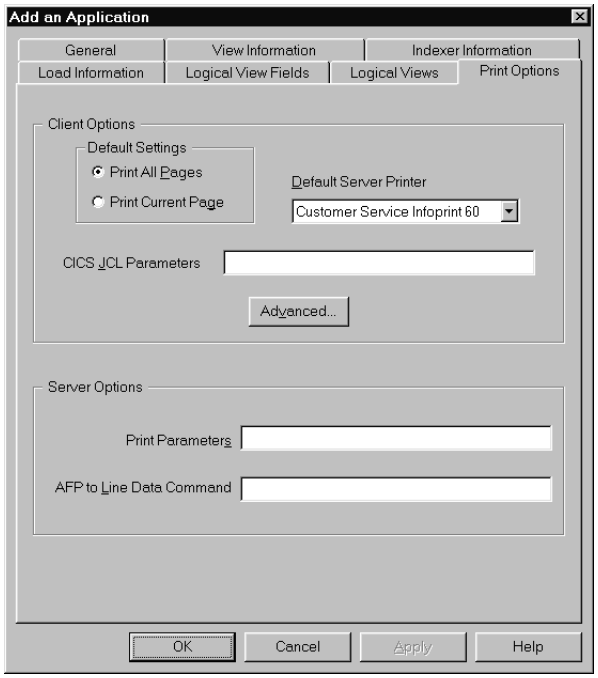


Figure 33. Print Options page

47. At this point, the properties of the application meet the requirements. If, however, you need to use an OS/400 printer file to further define some of the server print parameters, see Appendix C, "Server printing and faxing" on page 141 for details. Click OK in the Add an Application window. The OnDemand Administrator adds the application to the database and returns to the main window.

About the folder

Before we add the folder, let's review the data access requirements, the types of permissions that we need to specify, and the search and display fields that we need to define.

Data access requirements: The folder that we are defining allows users to access the telephone bill report application group and the telephone bill report application.

Permissions: Who needs access to the folder and what types of permissions do the users need?

- Users in the customer service department can open the folder to search for and retrieve statements.
- Users at Customer XYZ can open the folder to search for and retrieve statements that contain their account number and customer name.
- We will define a set of folder fields for the Customer XYZ users. The folder fields will limit access to specific statements.

Search and display fields: We need to define two sets of folder fields:

- One set at the folder level. These folder fields allow users in the customer service department to access any statement in the database.
- One set for the CustomerXYZ group. These folder fields allow users at Customer XYZ to access specific statements.

Adding the folder

Let's use the New Folder command to add the folder.

1. First, point to Folders and click the right mouse button. From the pop-up menu, select New Folder to open the Add a Folder dialog box.

   The pages of the Add a Folder dialog box organize information about the folder into sections. The tabs show which page you are on: General, Permissions, Field Definition, Field Information, and Field Mapping.
2. We'll start by completing the General page. (Search for *folder,general page* in the online help for details about the options that you can specify.) In the Name field, enter the name of the folder.
3. In the Description field, enter up to 120 characters of descriptive information about the folder.
4. Select the Display Document Location check box. This provides users with a visual clue about the type of media on which a statement is stored.
5. In the Application Groups list, select Telephone Bill Reports.

   Figure 34 shows the completed General page.

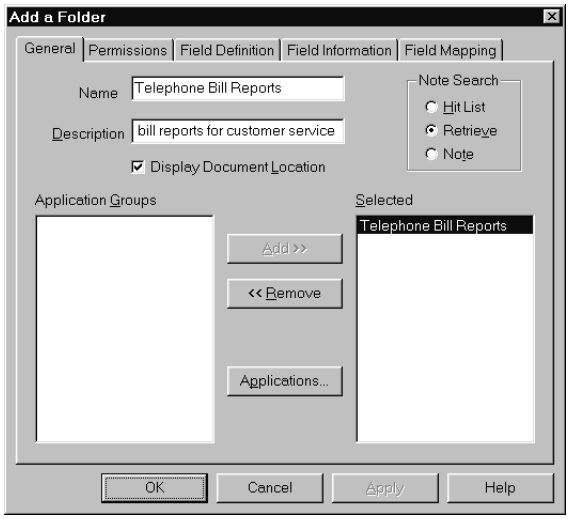


Figure 34. General page

6. Select the Permissions tab to specify the types of folder functions that users can perform. For example, we can let users open the folder, create private named queries, and maintain folder fields. We can specify default permissions and permissions for specific groups and users. By default, the person that creates the folder is given all folder permissions; no other users can open or maintain the folder. On our system, other users obtain permissions from a group. We need to add two groups to the folder. We defined the groups in "Adding the customer service group" on page 69 and "Adding the Customer XYZ group" on page 70.
7. From the Users and Group list, select +CustomerService.
8. Select the Access check box.

9. Click Add, to add the Customer Service group to the folder.
10. From the Users and Group list, select +CustomerXYZ.
11. Click Add, to add the CustomerXYZ group to the folder.

    Figure 35 shows the Permissions page. (Search for *folder,permissions* in the online help for details about the options that you can specify.)

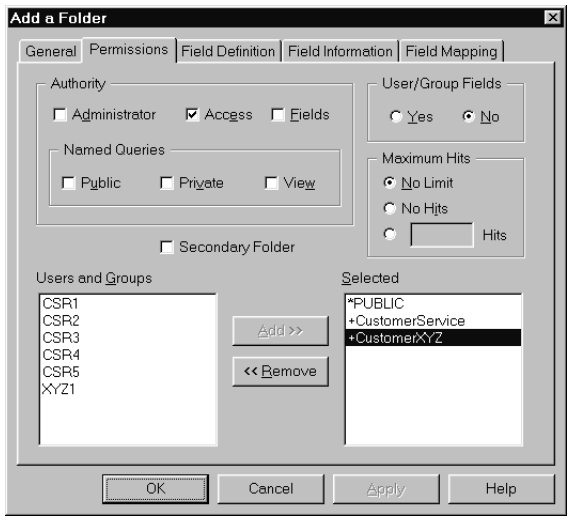


Figure 35. Permissions page

12. Select the Field Definition tab to define the folder fields. We need to define four folder fields to allow users to search for statements:
    - Account Number, a string field
    - Customer Name, a string field
    - Report Date, a date field
    - Other Information, a text search field
13. Complete the following steps to define a folder field:
    - In the Name field, enter the name of the folder field.
    - In the Description field, enter up to 120 characters of descriptive information about the folder field.
    - From the Field Type list, select the data type of the field. For the other Information field, select a data type of text search.
    - Select the Mapping Type. All of the fields in this example use the Single mapping type.
    - Click Add.

    Figure 36 on page 90 shows the completed Field Definition page. (Search for *folder,field definition page* in the online help for details about the options that you can specify.)

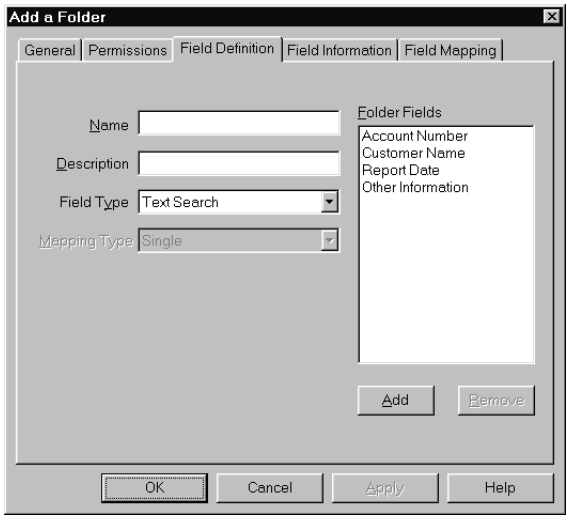


Figure 36. Field Definition page

14. Select the Field Information tab to specify the properties of the folder fields. Using the *PUBLIC identifier, you can specify field information that is used by all users that can open the folder. You can also specify field information for specific users and groups. The public field information will be used unless it is overridden by field information for a specific user or group. For this folder, we need to do the following:
    - Specify public field information. For the Account Number, Customer Name, and Other Information fields, we can accept the default values. For the Report Date field, we will specify field information.
    - Specify field information for the CustomerXYZ group. For the Other Information and Report Date fields, we can accept the default values. For the Account Number and Customer Name fields, we will specify field information.

    Search for *folder,field information page* in the online help for details about the options that you can specify.
15. First, specify the public field information for the Report Date field. See online help for *date format* or *time format* values for the Display Fmt and Defaults Fmt fields.
    - From the Name list, select Report Date
    - From the ID list, select *PUBLIC
    - Select the Default check box
    - From the Display Fmt list, select %b %e, %Y
    - From the Defaults Fmt list, select %b %e, %Y
    - In the Interval area, select Last, enter a 3 (three) in the entry field, and select Months

    Figure 37 on page 91 shows the Field Information page for the Report Date field.

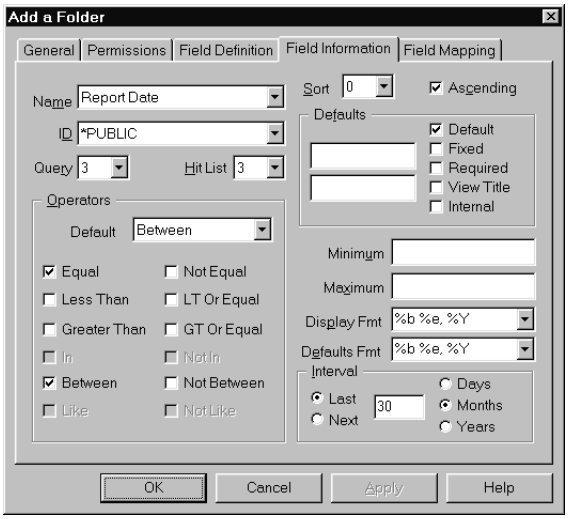


Figure 37. Field Information page

16. Next, specify the field information for the CustomerXYZ group. First we need to make a copy of the folder fields. Click the Permissions tab.
17. From the Selected List, select +CustomerXYZ.
18. In the User/Group Fields area, click Yes.

    Figure 38 shows the completed Permissions page.

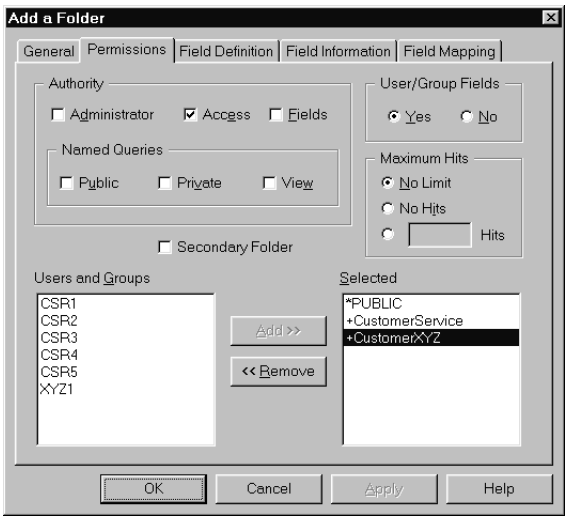


Figure 38. Permissions page

19. Click the Field Information Tab.
20. Specify the Customer XYZ field information for the Account Number field.
    - From the Name list, select Account Number
    - From the ID list, select +CustomerXYZ
    - From the Default list, select Equal
    - Clear the Like check box
    - Select the Default check box
    - Select the Fixed check box
    - In the first Defaults entry field, enter 1234567890123

- Clear the Append check box

Figure 39 shows the Field Information page for the Account Number field.

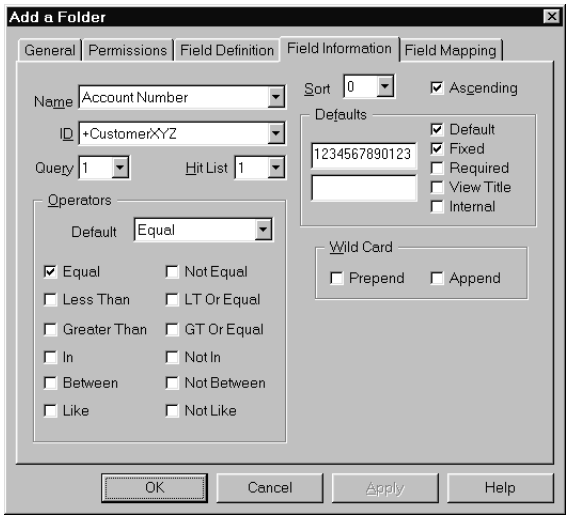


Figure 39. Field Information page

21. Next, specify the Customer XYZ field information for the Customer Name field.
    - From the Name list, select Customer Name
    - From the Default list, select Equal
    - Clear the Like check box
    - Select the Default check box
    - Select the Fixed check box
    - In the first Defaults entry field, enter Customer XYZ
    - Clear the Append check box

    Figure 40 on page 93 shows the Field Information page for the Customer Name field.

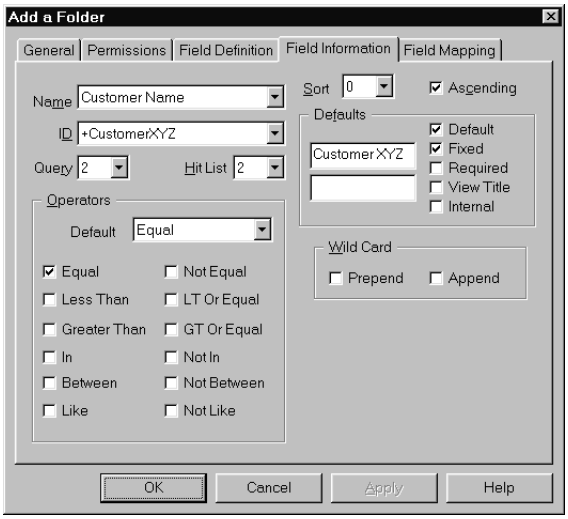


Figure 40. Field Information page

22. Select the Field Mapping tab to map the folder fields to application group database fields. OnDemand uses the values that users type in folder fields to construct SQL queries. An SQL query uses the database field name. Search for *folder,field mapping page* in the online help for details about the options that you can specify.
23. Map the following folder fields to their corresponding application group fields:
    - Account Number to acct
    - Customer Name to custname
    - Report Date to rdate

    Figure 41 shows the completed Field Mapping page.

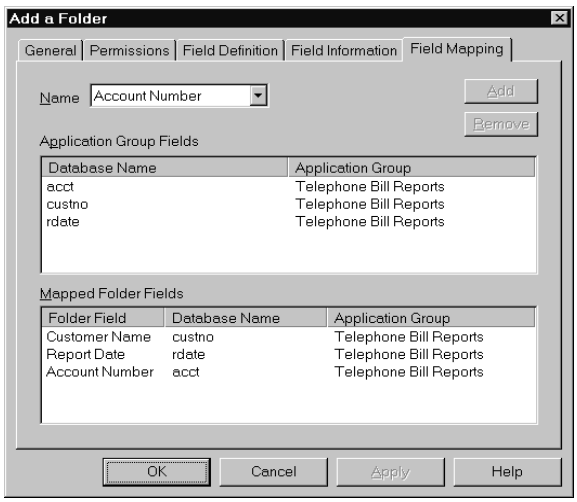


Figure 41. Field Mapping page

24. At this point, the properties of the folder meet the requirements. Click OK in the Add a Folder window. The administrative client adds the folder to the database and returns to the main window.

Summary

We have completed the basic requirements for adding a report to the system. We hope that the scenario that we described and developed is similar to how you plan to use OnDemand at your company. We have tried to add enough variations to show the flexibility of OnDemand to meet a range of business and operational requirements. Of course, there are several things we didn't show you. For example, we didn't create logical views of the report. Nor did we show you how to process the records that are saved in the system log. We also didn't show you every way to complete a given task or implement specific requirements. As with most administrative software, there is more than one way to accomplish something. Hopefully, we have shown you the most straightforward way to get things done with the administrative client.

You can use reference information provided with the product and the online help to find out more about how to use OnDemand. If you have questions and can't find the answers, please contact the IBM support center. We also offer classes that further explore how to administer the system.

Part 4. Loading data

This part of the book describes how to store data in OnDemand. The OnDemand data indexing and loading programs can process many types of input data, including AFP data, line data, unformatted ASCII data, Adobe PDF files, and other types of user-defined data. You can also load compressed and uncompressed, single and multi-page TIF files into the system. OnDemand can also process input files that contain graphics data, such as BMP, GIF, JFIF, and PCX. When you store data in OnDemand, the data indexing and loading programs extract index data from or generate index data about the input data. The index data is loaded into the OnDemand database. The input data is divided into indexed groups of pages (documents), compressed into storage objects, and copied to storage volumes. The OnDemand database is updated with the index information that allows users to search for and retrieve documents using OnDemand client programs. You can store copies of documents on disk and archive storage.

© Copyright IBM Corp. 2001, 2002

Chapter 10. Loading spooled file data

This section provides an overview of the data loading process — the process of adding the index data to the database and loading the report data and resources into OnDemand. When a load process completes, you can view the messages that were saved in the system log. To complete the data loading process, you should backup the OnDemand data on a regular basis.

Overview

You can archive the print output of your existing application programs in OnDemand without changing the print data stream or writing programs to process the data. In OnDemand, the print output of an application program is called a report. OnDemand provides programs that can index the reports, add the index data to the database, divide the input data into indexed groups of pages (documents), compress the documents, and copy the compressed documents into OnDemand. After you archive a report in OnDemand, your users can query, retrieve, and view or print pages of the report using the OnDemand client program.

The OnDemand data indexing and loading programs process input files that reside on the OnDemand server. If you generate your reports on another system, then you would typically transfer the reports to the OnDemand server and use the data indexing and loading programs to process them. When you index a report that contains AFP data, you must make sure that the data indexing program can access the resources required by the report. Resources include page segments and fonts. If the resources are not already on the OnDemand server, you must transfer them before loading the report.

You can create up to 32 index fields for each type of report that you define to OnDemand, providing many ways for users to query information contained in a report. The number of index fields that you define depends on the organization of the data in the report. For example, when you index a report that contains logical items, such as policies and statements, you might define index fields for the date, customer name, customer number, balance due, transaction number, and amount. When you index a report that contains transaction data, such as a general ledger, you might define index fields for the date and transaction number. After you determine what index fields you need and define them to the system, OnDemand extracts the index values from a report during the load process and stores them in records that are added to the database.

OnDemand compresses report data into storage objects, using information that you specify in the application. Depending on how you configure storage management for your application groups, OnDemand can automatically copy the report to disk and archive storage.

The load process saves messages in the system log each time that you load an input file into the system. You can open the System Log folder and view the messages for information such as the name of the input file, the indexing information, and the number of rows that were added to the database.

© Copyright IBM Corp. 2001, 2002

Preparing to load reports

There are a number of things to consider when preparing to load reports.

Storage space

When you initially configure an OnDemand system, you calculate the total amount of disk, optical, and tape storage required to hold the OnDemand database, database log files and reports. You also need to plan for the temporary space needed by OnDemand programs. The amount of storage space that you need on your system is usually a factor of how much data that you plan to store in OnDemand, how long you need OnDemand to maintain the data, the compression ratio that you can expect to achieve on the report data, and the number of copies of reports that you need the system to maintain. If you plan to index your reports on the OnDemand server, then you must allocate temporary space for the data indexing program. Temporary space is also required for the data loading program.

The *IBM Content Manager OnDemand for iSeries Common Server: Planning and Installation Guide* provides information, formulas, and worksheets that can help you calculate your storage requirements.

Defining the application group

When you archive reports in OnDemand, the ADDRPTOND command adds index data to the database and compresses report data into storage objects and copies the storage objects into OnDemand. The index data and the storage objects are associated with an application group. As part of defining a report to OnDemand, you must specify the application group that OnDemand uses to maintain the data. The application group specifies the database fields that hold the index data and the storage management information that determines where OnDemand maintains the report data and how long it maintains the index data and the report data. You can use the OnDemand administrative client to define an application group. See the online help for the administrative client for more information about defining database fields and specifying storage management information for application groups.

If you plan to maintain a copy of your reports in archive storage, then the application group must specify a migration policy and storage set that identify your storage management requirements.

Defining the application

Most customers define an application for each different report (or source of data) that they plan to archive in OnDemand. When you create an application, you specify information about the report, such as:

- The application group in which you want to store the report
- Physical information about the report, including the type of data found in the report
- The indexing parameters
- The processing that OnDemand should do to the index data before adding it to the database

You can use the OnDemand administrative client to define an application. See the online help for the administrative client for more information about defining applications.

Loading reports

You can load reports into OnDemand in a number of different ways, depending on your application requirements.

Running the ADDRPTOND command

The ADDRPTOND command is the primary OnDemand data indexing and loading command. The ADDRPTOND command determines if the input data needs to be indexed, and if it does, calls the indexing program. The ADDRPTOND command then processes the index data, adding it to the database, optionally compresses the report data into storage objects, and copies the storage objects to storage volumes.

You run the ADDRPTOND command each time that you want to load a report or set of reports into the system. You can either run the command from the command line or use the OnDemand output queue monitor to periodically check for input data to process. See "Output queue monitors" below for more detail.

Note: The OnDemand server program must be running, otherwise the ADDRPTOND command will fail.

Output queue monitors

You can use the Start Monitor for OnDemand (STRMONOND) command to start a monitor program for any output queue to receive spooled files for archiving. The monitor continuously checks an output queue for spooled files, and allows OnDemand to capture the spooled files as they arrive.

We recommend using the monitors in a batch environment only.

You generally define and then manually store reports with the ADDRPTOND command during testing. Then, when the report is ready for production, you can automate storage with the OnDemand monitor—if you use the *SPLFNAME, *FORMTYPE, *USERDATA, *JOBNAME, or *USRDFNxxxx spooled file attribute of the report to match the application group and application names you defined to OnDemand.

The spooled files must be in ready status (RDY) on the output queue that you are monitoring. The first time a Monitor is started for an output queue, OnDemand will create a data queue and attach it to the output queue being monitored. If there are spooled files already in that queue, the monitor will not detect them. If this happens, place these spooled files on hold, then release them. The Monitor should then detect and process them. This should only be necessary the first time a Monitor job is started for a particular output queue.

To begin capturing reports automatically when a spooled file arrives in a particular output queue, you can issue the STRMONRDAR command. The monitor runs continuously until the End time or Number of hours occurs. You can also use the End Monitor for OnDemand (ENDMONOND) command to stop processing.

The monitor processes any spooled files that arrive in the named output queue in RDY (Ready) status. Spooled files with any other status are not affected by the monitor, and remain in the output queue until they are deleted or moved.

You may wish to add the STRMONOND command to your system startup program so the monitor(s) start each time you IPL the system. See the online help for more information about these commands and their parameters.

Note: The OnDemand server program must be running, otherwise the spooled files will fail to load.

Indexing reports

You must generate index data for a report before you load the report into OnDemand. If the report contains AFP data, then you need to store the AFP resources in OnDemand. The resources are required to display and reprint pages of a report that contains AFP data. There are several programs provided with OnDemand to help you generate index data for your reports:

- OnDemand OS/400 Indexer, to index most OS/400 spooled files
- OnDemand PDF Indexer, to index Adobe PDF files
- OnDemand Generic Indexer, to index almost any other type of input data

For details about indexing data, including how to use other indexers, see the *IBM Content Manager OnDemand for iSeries Common Server: Indexing Reference*.

Processing the input data

After a report has been indexed, the ADDRPTOND command can process the index file to prepare the index data for loading into the database and prepare the report data and resource group files for the storage managers to load on storage volumes.

Processing index data

The index data is processed by the ADDRPTOND command before the index data is added to the database. The ADDRPTOND command extracts information from the application group and the application and performs the following processing:

- Determines the database field information from the application group.
- Determines the preprocessing information from the application, if needed.
- After preprocessing the index file, the ADDRPTOND command creates the database rows:
  - One row for every group of indexed pages in a report that contains a sorted transaction value
  - One row for every indexed item in a report that contains logical items, such as policies and statements
- Determines the postprocessing information from the application. For example, the ADDRPTOND command may need to drop duplicate index records. You specify exactly what processing you want the command to do when you define the application.
- Passes the index rows to the database manager. The rows consist of fields that contain the index values that the indexing program extracted from the report and other fields generated by OnDemand. An index row contains:
  - One column for each field defined in the application group
  - One or more columns of OnDemand control information

Processing reports and resources

The ADDRPTOND command divides the input data into indexed groups of pages (documents) and compresses the documents into storage objects. Dividing a report into groups of pages improves the efficiency of queries and can improve the time required to retrieve and display the report. Compression improves the efficiency of the storage manager. In this step, the ADDRPTOND command:

- Determines the compression information from the application.

- Determines the storage management information from the application group. The storage management information determines the storage locations, such as disk storage and archive storage.
- Locates the AFP resources called for in the input data. If the resources are not available, the load will fail. Note that AFP overlays are supported only when using the AFP data type. If you have SCS spooled files with AFP overlays and you need to keep the overlay with the print data for viewing, printing, and so forth, then your printer file needs to be changed to *AFPDS prior to creating the spooled file. Then, when defining the report to OnDemand, the application's Data Type field would be set to AFP.
- Compresses the documents into storage objects. OnDemand compresses report data into approximately 100 KB blocks (a default value) and places the blocks into a storage object. OnDemand uses a 10 MB storage object (a default value) to improve storage efficiency and performance. Compressed data does not span storage objects. OnDemand assigns unique file names to the storage objects and sequentially numbers them within an application group.

Loading index data

The ADDRPTOND command works with the database manager to add the index data to the database. Depending on the database organization that you specified when you defined the application group, OnDemand either creates a new table each time that you load a report or adds the index data to an existing database table.

OnDemand uses a *segment table* as a high level index to the index data for an application group. Each row in the segment table identifies a specific table of application group index data. The fields in the segment table identify the application group and the dates found in a table. The dates represent the earliest and latest dates that can be found in that segment of application group index data. OnDemand can use the segment table to limit a query to a specific table of application group index data.

The database manager updates the segment table if the beginning date in the report is earlier than a date already stored in the table.

OnDemand limits the size of a table to improve performance and storage management. The number of rows in a table can be specified when you define the application group. The default size of a table is ten million rows. The database manager automatically closes a table and opens a new table when this threshold is reached. When closing a table, the database manager updates the segment table with the latest ending date found in the table.

Loading storage objects

The ADDRPTOND command calls the storage manager to copy storage objects to disk and archive storage. The storage manager extracts information from the application group to determine where and when to copy the storage objects.

OnDemand uses an object called a storage set to determine the locations that can hold report data. A storage set and its associated migration policy (of the exact same name) point to one or more storage locations. A storage set can write data to one and only one storage location at a time (the *active* storage level).

Disk storage

The primary purpose of disk storage is for short-term, high-speed retrieval of report data.

If you configure your application groups to copy data to disk storage, then the disk storage manager copies the storage object to disk. The `Cache Data for xx Days` setting on the Storage Management page determines whether OnDemand copies documents to disk storage.

Archive storage

A storage set can identify an archive storage media such as optical or tape. OnDemand uses its archive storage manager to maintain storage objects in archive storage for long-term storage and for backup copies of reports.

The storage manager can copy the storage object to archive storage when the report is initially loaded into the system or at a later time, depending on how you configure your application groups. Most customers configure the system to copy report data to disk and archive storage at the same time.

Resources

OnDemand always stores resources (such as AFP overlays or page segments) on disk, to provide fast retrieval when a user selects an item for viewing. OnDemand saves only one copy of a resource on the system, even if several reports use the same resource. When processing a resource group file, the ADDRPTOND command checks the resource identifier to determine if the resource is already present on the system.

Verifying processing

When you add a report into the system, the ADDRPTOND command saves a copy of the messages generated during the load process in the system log. After a load process completes, you can open the System Log folder and view the messages. The messages reference the load program named ARSLOAD and include the date and time that the load process started and completed, the name of the input file, and the number of rows that were added to the database. For example:

```
arsload: Processing file >br1010a<
arsload: Fri Mar 29 12:37:23 MST 1996 Indexing started,
          394419060 bytes to process
arsload: Fri Mar 29 12:41:33 MST 1996 Indexing completed
arsload: Fri Mar 29 12:41:35 MST 1996 Load Started,
          394419060 bytes to load
  Resource br1010a.res matches the resource 2
  OnDemand Load ID = >6850-25-0-15FAA-9577-9577<
  Loaded 130000 rows into the database
arsload: Fri Mar 29 13:05:14 MST 1996 Load Completed
arsload: Processing successful for file >br1010a<
```

You can verify the number of rows that OnDemand added to the database:

- For a report that contains transaction data that you have divided into indexed groups of pages, the number of pages in the report divided by the number of pages in an indexed group of pages should equal the number of rows added to the database. For example, if a report contains 150,010 pages and there are 100 pages in an indexed group of pages, then OnDemand should have added 1,501 rows to the database.
- For a report that contains logical items, such as statements and policies, the number of rows added to the database should equal the number of indexed items in the report. For example, if a report contains 1,000 statements, OnDemand should have added 1,000 rows to the database.

The OnDemand Load ID represents the data that the ADDRPTOND command stored in the system during a load process. The Load ID can be used to identify a specific load process. For example, you can run the RMVRPTOND command and

specify the Load ID to delete the index data and documents that were created when the ADDRPTOND command processed an input file. Table 5 lists the fields in a Load ID.

Table 5. OnDemand Load ID

Load ID Field	Meaning
6850	Application group identifier
25	Primary node identifier
0	Secondary node identifier
15FAA	The load identifier within the application group
9577	The earliest date in the report (relative to 1/1/1970)
9577	The latest date in the report (relative to 1/1/1970)

Backing up databases

After you load reports into OnDemand, we recommend that you create a backup copy of the OnDemand data.

- When you backup the OnDemand database, you protect control information and index data that OnDemand and the database manager need to support the system.
- When you backup the OnDemand data in IFS, you protect your data that resides on disk that may not have been copied to archive media yet.

We recommend that you backup the data at least once a week, and more often if you load reports every day.

See Chapter 5, “Backup and recovery” on page 21 for details about OnDemand backup and recovery.

Chapter 11. Loading image files

Overview

OnDemand provides support for storing and retrieving letters and other types of correspondence. The most straight forward way to store this type of information is to transform a printed copy of a letter into an image file using a scanner and image capture software. In addition to scanning the letter and saving it as an image file, you must create index data for the letter. You can then use the ARSLOAD program to load the index data into the database and store the letter file on the system. Your users can then query, retrieve, and view, print or send copies of the letter using one of the OnDemand client programs.

OnDemand provides the Generic Indexer so that you can index input files that contain data other than AFP data, line data, and PDF. You specify the index data for the input files that you want to index with the Generic Indexer in a parameter file. The parameter file contains the index field names and values and identifies the input files that you want to process. You can create up to 32 index fields for each input file that you want to load into the system, providing many ways for users to query and retrieve the letters. The number of index fields that you define usually depends on how your users retrieve documents from the system. In the example that follows, we defined index fields for the date on the letter, the name of the person that sent the letter, the company name, and the subject of the letter.

This section describes how to use a scanner and image capture software to create the image files. Before you can load the image files into the system, you must define an application group to manage the storage of the letters, an application to specify the physical and logical properties of the letters, and a folder to let users search for and retrieve the letters.

Defining the application group

When you load an input file into the system, OnDemand updates the database with the index data that you provide and stores the indexed groups of pages as documents in cache storage and archive storage. The application group contains the information that OnDemand uses to store and maintain the index data and the documents on the system. This section provides information about some of the key properties of the application group. You can use the OnDemand administrative client to define an application group.

Database Organization

The Database Organization determines how OnDemand organizes the index data that is stored in the application group.

We accepted the default Database Organization of Multiple Loads per Database Table and Table Size of its maximum number of rows. Each time that we load a letter (input file) into the system, OnDemand adds one row to a database table. When a table reaches its maximum number of rows, OnDemand closes the table and creates a new table. OnDemand always adds index records to the open table; the closed tables can be queried.

© Copyright IBM Corp. 2001, 2002

Expiration Type

The Expiration Type determines how OnDemand deletes index data and documents from the application group.

We accepted the default Expiration Type of Load. This means that OnDemand deletes the index data and documents from one load process (one or more input files) at a time from the application group. Depending on the number of input files processed during a load process, one or more letters (input files) may be deleted at a time. For example, if we create a parameter file for the Generic Indexer that contains index data for several input files, then OnDemand will delete all of the index data and documents from the application group at the same time.

Permissions

Application group permissions determine the users that can access information stored in the application group and determine the users that can do other types of tasks related to the application group.

Under the *PUBLIC identifier, we specified the Access permission so that all of the users defined to the server can access data stored in the application group.

Field Definition

The Field Definition page is where you define the database fields for the application group. When you load an input file into the system, OnDemand stores the index data that is specified in the parameter file into fields in records that are added to the database. When a user searches for letters, OnDemand compares the search criteria entered by the user with index data in the application group.

We defined the following database fields. The fields allow users to locate letters based on different criteria, such as the date of the letter, the name of the person that sent the letter, and the subject of the letter.

ldate The date on the letter. Defined as a date field.

name The person that sent the letter. Defined as a string field that contains variable length data.

company The person's company. Defined as a string field that contains variable length data.

subject The subject of the letter. Defined as a string field that contains variable length data.

Defining the application

Most customers define an application for each different source of input data that they plan to load into OnDemand. This section provides information about some of the key properties of the application. You can use the OnDemand administrative client to define an application.

Application Group

You must assign an application to an application group. We assigned the application to the application group that we created in "Defining the application group" on page 105.

Data Format

The Data Format determines the kind of data that is found in the documents that are stored on the system. If the type of the data that is found in the original input file is different than the type of the data that is found in the documents that are stored on the system, then you should specify the type of the data that will be stored in OnDemand. (In this example, the type of the data is the same – TIFF.) We selected TIFF as the Data Type. When we save the scanned image of a letter, we must save the file as a TIF file.

Indexer

The Indexer determines the indexing program that OnDemand uses to index and convert input data. We selected Generic as the Indexer. To store TIF files on the system, we must index them with the Generic Indexer program provided with OnDemand.

Data Compression

The Data Compression determines whether OnDemand compresses the input files.

We selected None, so that OnDemand does not compress the input files. Since we plan to save the scanned image files as compressed TIF files, we do not need OnDemand to compress them. If you need to save scanned image files as uncompressed TIF files, then you should select one of the compression methods supported by OnDemand. That way, OnDemand can compress the input files to improve storage efficiency and retrieval performance.

Defining the folder

You must define a folder so that users can search for and retrieve the input files that you load into OnDemand. This section provides information about some of the key properties of the folder. You can use the OnDemand administrative client to define a folder.

Application Group

A folder can be used to search one or more application groups. We selected the application group that we created in “Defining the application group” on page 105. When users open the folder, they can search for and retrieve the input files that we loaded into the application group.

Permissions

Folder permissions determine the users that can open the folder and determine the users that can do other types of tasks related to the folder.

Under the *PUBLIC identifier, we specified the Access permission so that all users defined to the server can open the folder.

Field Definition

The Field Definition page is where you define the search fields for the folder. The search fields contain the search criteria entered by the user. For most folders, you probably want to define a search field for each field that you defined for the application group.

We defined the following fields. The fields allow users to locate letters based on different criteria, such as the date of the letter, the name of the person that sent the letter, and the subject of the letter.

Letter Date The date on the letter.

From The person that sent the letter.

Company The person's company.

Subject The subject of the letter.

Field Mapping

The Field Mapping page is where you *map*, or associate, the folder fields to the application group fields. This is how you specify that the search criteria that a user enters in a particular folder field should be used to search a specific application group field. We mapped each of the folder fields to their corresponding application group fields. For example, we mapped the folder field named `Letter Date` to the application group field named `Ldate`.

Acquiring the images

We scanned each letter using an IBM 3119 PageScanner. We previewed and saved the image files using the IBM Image Capture Facility. Each letter was saved to a separate file as a compressed TIF file.

We specified the following settings to the image capture software:

- Paper Size 8.5 inches
- Bi-level image
- 300 DPI
- Compressed TIFF

If you plan to let OnDemand compress the data, then you should save the files as uncompressed TIF files.

Accessing the image files

The ARSLOAD program runs on the OnDemand server, and is the primary OnDemand program for loading images and other user-defined data. The files that you want to load into the system must be stored on the server or you must provide network access to the files.

You can transfer the files to the server using one of a number of methods. We used FTP to copy the files from the PC that ran the scanning operation to the server where we plan to run the ARSLOAD program. We copied the image files from the PC to the server as binary files with a file type of `TIF`.

Creating index data

You must create index data for your input files before you can load them into OnDemand. If the input data is other than AFP, line, or PDF, then you must use the Generic Indexer to load it into the system. You specify index data for the Generic Indexer using a parameter file. You should use a standard text editor to create the parameter file. Figure 42 on page 109 shows an example of a parameter file that can be used by the Generic Indexer to process image files. See the *IBM*

Content Manager OnDemand for iSeries Common Server: Indexing Reference for more information about the Generic Indexer and the parameter file.

```
COMMENT:          input file number 1
GROUP_FIELD_NAME:ldate
GROUP_FIELD_VALUE:09/01/95
GROUP_FIELD_NAME:name
GROUP_FIELD_VALUE:Mr. Earl Hawkins
GROUP_FIELD_NAME:company
GROUP_FIELD_VALUE:Soft Products
GROUP_FIELD_NAME:subject
GROUP_FIELD_VALUE:optical storage devices
GROUP_OFFSET:0
GROUP_LENGTH:0
GROUP_FILENAME:letter1.tif
COMMENT:          input file number 2
GROUP_FIELD_NAME:ldate
GROUP_FIELD_VALUE:09/01/95
GROUP_FIELD_NAME:name
GROUP_FIELD_VALUE:Hans G. Piker
GROUP_FIELD_NAME:company
GROUP_FIELD_VALUE:MBI Company
GROUP_FIELD_NAME:subject
GROUP_FIELD_VALUE:optical storage devices
GROUP_OFFSET:0
GROUP_LENGTH:0
GROUP_FILENAME:letter2.tif
COMMENT:          input file number 3
GROUP_FIELD_NAME:ldate
GROUP_FIELD_VALUE:09/16/95
GROUP_FIELD_NAME:name
GROUP_FIELD_VALUE:Laurie Unicolas
GROUP_FIELD_NAME:company
GROUP_FIELD_VALUE:Dove Properties
GROUP_FIELD_NAME:subject
GROUP_FIELD_VALUE:account balance due
GROUP_OFFSET:0
GROUP_LENGTH:0
GROUP_FILENAME:letter3.tif
COMMENT:          input file number 4
GROUP_FIELD_NAME:ldate
GROUP_FIELD_VALUE:10/01/95
GROUP_FIELD_NAME:name
GROUP_FIELD_VALUE:George VanLocal
GROUP_FIELD_NAME:company
GROUP_FIELD_VALUE:Express American
GROUP_FIELD_NAME:subject
GROUP_FIELD_VALUE:airline fares
GROUP_OFFSET:0
GROUP_LENGTH:0
GROUP_FILENAME:letter4.tif
```

Figure 42. Example of a Generic Indexer parameter file

We saved the parameter file as `LETTERS.IND` on the server in the directory from which we plan to run the ARSLOAD program.

Configuring the ARSLOAD program

Note: The OnDemand server program must be running, otherwise the ARSLOAD program will fail.

The ARSLOAD program is the primary OnDemand data indexing and loading program for image files. The ARSLOAD program determines if the input data needs to be indexed, and if so, it calls the indexing program. The ARSLOAD program then processes the index data, loading it into the database, optionally compresses the input data into storage objects, and copies the storage objects to storage volumes. See "ARSLOAD" on page 171 for more information about running the ARSLOAD program within QSHELL, including the parameters that you can specify to process input files.

You typically run the ARSLOAD program each time that you want to load a file or set of files into the system. You can either run the ARSLOAD program from the command line or configure it to run periodically to check for input data to process. To support a low volume scanning operation, we chose to run the ARSLOAD program from the command line using QSHELL. For example:

```
arsload -n -g Letters letters
```

We specified the following parameters to the ARSLOAD program:

-n Do not delete the input files.

-g Letters The name of the application group to load.

letters The name of the input file to process.

In the example, the ARSLOAD program can locate the input file `LETTERS.IND` in the current directory. The input file contains the index information that we created in "Creating index data" on page 108 (which is a parameter file for the Generic Indexer). The image files that we want to load are in the same directory.

Processing the input data

The ARSLOAD program processes the parameter file, loads the index data into the database, and loads the image files into OnDemand.

Processing index data

The ARSLOAD program processes the parameters that we specified on the command line and the indexing parameter file before loading the index data into the database. The ARSLOAD program extracts information from the application group and the application and performs the following processing:

- Determines the database field information from the application group
- Determines the data type and other information from the application
- Creates one database row for each group that we specified in the parameter file. We specified one group for each image file.
- Passes the index rows to the database manager. The rows contain the index values and other fields generated by OnDemand. An index row contains:
  - One column for each field that we defined for the application group
  - One or more columns of OnDemand control information

Processing the image files

The ARSLOAD program processes each image file that we specified in the parameter file. The ARSLOAD program extracts information from the application group and the application and performs the following processing:

- Determines the compression information from the application

- Determines the storage management information from the application group. The storage management information determines the storage locations, such as disk and archive storage.
- Stores the image files in storage objects. OnDemand uses a 10 MB storage object (the default value) to improve storage efficiency and performance. OnDemand assigns unique names to the storage objects and sequentially numbers them within an application group.

Verifying processing

The ARSLOAD program saves a copy of the messages that were generated during a load process in the system log. After a load process completes, you can open the System Log folder and review the messages. The information in the messages includes the date and time that the load process started and completed, the name of the input file(s), and the number of rows that were added to the database. For example:

```
arsload: Processing file >letters<
arsload: Mon Dec 18 13:04:37 MST 1995 Indexing started,
          116017 bytes to process
  OnDemand Load ID = >6927-0-0-1FAA-0-0<
  Loaded 4 rows into the database
arsload: Mon Dec 18 13:05:14 MST 1995 Loading Completed
arsload: Processing successful for file >letters<
```

You can verify the number of rows that OnDemand added to the database. In the example, the number of rows added to the database should equal the number of groups (and image files) that we specified in the parameter file.

The OnDemand Load ID represents the data that the ARSLOAD program stored into the system during a load process. The Load ID can be used to identify a specific load process. For example, you can run the RMVRPTOND command and specify the Load ID to delete the index data and documents that were created when the ARSLOAD program processed a Generic Indexer parameter file.

Chapter 12. Loading user-defined data

OnDemand provides support for storing and retrieving almost any type of data. For example, OnDemand provides support for AFP, Line, PDF, GIF, JFIF (JPEG), PCX, and TIFF data. However, OnDemand is not limited to maintaining these types of data. OnDemand provides the User-Defined data type to support almost any other type of data that you want to store in the system. For example, you can configure the system to process Lotus® WordPro documents, so that when a user retrieves one of the files from the system, OnDemand automatically starts Lotus WordPro to open the document.

To store user-defined data on the system, you must create index data for the input files and you must register the file type of the input file with OnDemand. The file type determines the program that is started to open a file when a user retrieves one of the files from the system. The file type must also be registered with the client operating system. If your OnDemand system supports client programs that run under different operating systems, then you must register the specified file type on all of the client operating systems.

OnDemand provides the Generic Indexer so that you can index user-defined data. You specify the index data for the input files that you want to index with the Generic Indexer in a parameter file. The parameter file contains the index field names and values and identifies the input files that you want to process. You can create up to 32 index fields for each input file that you want to load into the system, providing many ways for users to query and retrieve documents. The number of index fields that you define usually depends on how your users retrieve documents from the system. For example, you might want to define index fields for the date, author, and subject or purpose of the user-defined data.

Before you can load user-defined data into the system, you must define an application group to manage the storage of the files, an application to specify the physical and logical attributes of the input files, and a folder to let users search for and retrieve the files. This section contains an example that shows how to define Lotus WordPro files to OnDemand. It provides an overview of defining the application group, application, and folder. For details and a comprehensive example of defining input data to OnDemand, see Part 3, “Using iSeries Navigator for OnDemand administration” on page 23.

Defining the application group

When you load an input file into the system, OnDemand updates the database with the index data that you provide and stores the indexed groups of pages as documents in cache storage and archive storage. The application group contains the information that OnDemand uses to store and maintain the index data and the documents on the system. This section provides information about some of the key properties of the application group. You can use the OnDemand administrative client to define an application group.

Database organization

The Database Organization determines how OnDemand organizes the index data that is stored in the application group.

© Copyright IBM Corp. 2001, 2002

We accepted the default Database Organization of Multiple Loads per Database Table and Table Size of its maximum number of rows. Each time that we load a Lotus WordPro file into the system, OnDemand adds one row to a database table. When a table reaches its maximum number of rows, OnDemand closes the table and creates a new table. OnDemand always adds index records to the open table; closed tables can be queried.

Expiration type

The Expiration Type determines how OnDemand deletes index data and documents from the application group.

We accepted the default Expiration Type of Load. This means that OnDemand deletes the index data and documents from one load process (one or more input files) at a time from the application group. Depending on the number of input files processed during a load process, one or more Lotus WordPro files may be deleted at a time. For example, if we create a parameter file for the Generic Indexer that contains index data for several input files, then OnDemand will delete all of the index data and documents from the application group at the same time.

Permissions

Application group permissions determine the users that can access information stored in the application group and determine the users that can do other types of tasks related to the application group.

Under the *PUBLIC identifier, we specified the Access permission so that all of the users defined to the server can access data stored in the application group.

Field definition

The Field Definition page is where you define the database fields for the application group. When you load an input file into the system, OnDemand stores the index data that is specified in the parameter file into fields in records that are added to the database. When a user queries the system, OnDemand compares the search criteria entered by the user with index data in the application group.

We defined the following database fields. The fields allow users to locate files based on different criteria, such as the date, author, and subject of the file.

fdate The date associated with the file. For example, the date the file was created or the date the file was published. Defined as a date field.

author The author of the file. Defined as a string field that contains variable length data.

subject The subject or purpose of the file. Defined as a string field that contains variable length data.

Defining the application

Most customers define an application for each different source of input data that they plan to load into OnDemand. This section provides information about some of the key properties of the application. You can use the OnDemand administrative client to define an application.

Application Group

You must assign the application to an application group. We assigned the application to the application group that we created in "Defining the application group" on page 113.

Data Format

The Data Format determines the kind of data that is found in the documents that are stored on the system. If the type of the data that is found in the original input file is different than the type of the data that is found in the documents that are stored on the system, then you should specify the type of the data that will be stored in OnDemand.

Because there is not a supplied Data Format for Lotus WordPro files, we must select User Defined from the Data Type list.

File Extension

When you select User Defined from the Data Type list, you must also enter a value in the File Extension. The File Extension determines the program that is started to open a user-defined file when it is retrieved from the system. In our example, we entered the characters `LWP`, for Lotus WordPro.

The File Extension that you specify must also be registered to the client operating system. For example, in Windows NT, start Explorer. From the View menu, select Options. The File Types page shows the file types that are currently registered with the system. Verify that the file type that you specified in the File Extension field is registered with the system. If it is not, then you must define a new type. See the online help for details. If your OnDemand system supports client programs that run under different operating systems, then you must register the file type on all of the client operating systems.

Indexer

The Indexer determines the indexing program that OnDemand uses to index and convert input data. We selected Generic as the Indexer. To store user-defined files on the system, we must index them with the Generic Indexer program provided with OnDemand.

Data Compression

The Data Compression determines whether OnDemand compresses the input files.

On the Load Information page, we can accept the default Data Compression of `OD77`, because we want OnDemand to compress input files before storing them on the server. Resources are not supported with user-defined input data.

Defining the folder

You must define a folder so that users can search for and retrieve the input files that you load into OnDemand. This section provides information about some of the key properties of the folder. You can use the OnDemand administrative client to define a folder.

Application Group

A folder can be used to search one or more application groups. We selected the application group that we created in "Defining the application group" on page 113.

When users open the folder, they can search for and retrieve the Lotus WordPro files that we loaded into the application group.

Permissions

Folder permissions determine the users that can open the folder and determine the users that can do other types of tasks related to the folder.

Under the *PUBLIC identifier, we specified the Access permission so that all users defined to the server can open the folder.

Folder fields

The Field Definition page is where you define the search fields for the folder. The search fields contain the search criteria entered by the user. For most folders, you probably want to define a search field for each field that you defined for the application group.

We defined the following folder fields. The fields allow users to locate files based on different criteria, such as the date, the author, and the subject.

File Date The date associated with the file. For example, the date the file was created or the date that the file was published.

Author The person that created the file.

Subject The subject or purpose of the file.

Field Mapping

The Field Mapping page is where you *map*, or associate, the folder fields to the application group fields. This is how you specify that the search criteria that a user enters in a particular folder field should be used to search a specific application group field. We mapped each of the folder fields to their corresponding application group fields. For example, we mapped the folder field named `File Date` to the application group field named `fdate`.

Accessing the input files

The ARSLOAD program runs on the OnDemand server and is the primary data loading program for user-defined data. The files that you want to load into the system must be stored on the server or you must provide network access to the files, and is the primary OnDemand program to load user-defined and image data.

You can transfer the files to the server using one of a number of methods. We used FTP to copy the files from the PC on which the files were created to the server where we plan to run the ARSLOAD program. We copied the input files from the PC to the server as binary files with a file type of `LWP`.

Creating the index data

You must create index data for your input files before you can load them into the system. If the input data is other than AFP, line, or PDF, then you must use the Generic Indexer to load it into the system. You must specify index data for the Generic Indexer using a parameter file. You should use a standard text editor to create the parameter file. Figure 43 on page 117 shows an example of a parameter file that can be used by the Generic Indexer to process the Lotus WordPro files. See

the *IBM Content Manager OnDemand for iSeries Common Server: Indexing Reference* for more information about the Generic Indexer and the parameter file.

```
COMMENT:          input file number 1
GROUP_FIELD_NAME:fdate
GROUP_FIELD_VALUE:12/18/95
GROUP_FIELD_NAME:author
GROUP_FIELD_VALUE:Jessica Hawkins
GROUP_FIELD_NAME:subject
GROUP_FIELD_VALUE:optical storage devices
GROUP_OFFSET:0
GROUP_LENGTH:0
GROUP_FILENAME:letter1.lwp
COMMENT:          input file number 2
GROUP_FIELD_NAME:fdate
GROUP_FIELD_VALUE:12/18/95
GROUP_FIELD_NAME:author
GROUP_FIELD_VALUE:Paul Garveys
GROUP_FIELD_NAME:subject
GROUP_FIELD_VALUE:optical storage devices
GROUP_OFFSET:0
GROUP_LENGTH:0
GROUP_FILENAME:letter2.lwp
COMMENT:          input file number 3
GROUP_FIELD_NAME:fdate
GROUP_FIELD_VALUE:12/18/95
GROUP_FIELD_NAME:author
GROUP_FIELD_VALUE:Randy Perkinsen
GROUP_FIELD_NAME:subject
GROUP_FIELD_VALUE:account balance due
GROUP_OFFSET:0
GROUP_LENGTH:0
GROUP_FILENAME:letter3.lwp
COMMENT:          input file number 4
GROUP_FIELD_NAME:fdate
GROUP_FIELD_VALUE:12/18/95
GROUP_FIELD_NAME:author
GROUP_FIELD_VALUE:Georgia July
GROUP_FIELD_NAME:subject
GROUP_FIELD_VALUE:airline fairs
GROUP_OFFSET:0
GROUP_LENGTH:0
GROUP_FILENAME:letter4.lwp
```

Figure 43. Example of a Generic Indexer parameter file

We saved the parameter file as `LWP.IND` on the server in the directory from which we plan to run the ARSLOAD program.

Configuring the ARSLOAD program

Note: The OnDemand server program must be running, otherwise the ARSLOAD program will fail.

The ARSLOAD program is the primary OnDemand data indexing and loading program for user-defined data. The ARSLOAD program determines if the input data needs to be indexed, and if so, it calls the indexing program. The ARSLOAD program then processes the index data, loading it into the database, optionally compresses the input data into storage objects, and copies the storage objects to

storage volumes. See “ARSLOAD” on page 171 for more information about running the ARSLOAD program within QSHELL, including the parameters that you can specify to process input data.

You typically run the ARSLOAD program each time that you want to load a file or set of files into the system. You can either run the ARSLOAD program from the command line or configure it to run as a daemon (UNIX® servers) or service (Windows servers) to periodically check for input data to process. To store a few WordPro files at a time, we chose to run the ARSLOAD program from the command line using QSHELL. For example:

```
arsload -n -g 'Lotus WordPro Documents' lwp
```

We specified the following parameters to the ARSLOAD program:

-n Do not delete the input files.

-g 'Lotus WordPro Documents'
The name of the application group to load.

lwp The name of the input file to process.

In the example, the ARSLOAD program can locate the input file `LWP.IND` in the current directory. The input file contains the index information that we created in “Creating the index data” on page 116 (which is a parameter file for the Generic Indexer). Because we did not specify the full path name of the input files, they must reside in the same directory.

Processing the input data

The ARSLOAD program processes the parameter file, loads the index data into the database, and loads the Lotus WordPro files into OnDemand.

Processing index data

The ARSLOAD program processes the parameters that we specified on the command line and the indexing parameter file before loading the index data into the database. The ARSLOAD program extracts information from the application group and the application and performs the following processing:

- Determines the database field information from the application group
- Determines the data type and other information from the application
- Creates one database row for each group that we specified in the parameter file. We specified one group for each Lotus WordPro file.
- Passes the index rows to the database manager. The rows contain the index values and other fields generated by OnDemand. An index row contains:
  - One column for each field that we defined for the application group
  - One or more columns of OnDemand control information

Processing the Lotus WordPro files

The ARSLOAD program processes each Lotus WordPro file that we specified in the parameter file. The ARSLOAD program extracts information from the application group and the application and performs the following processing:

- Determines the compression information from the application
- Determines the storage management information from the application group. The storage management information determines the storage locations, such as disk and archive storage.

- Stores the Lotus WordPro files in storage objects. OnDemand uses a 10 MB storage object (the default value) to improve storage efficiency and performance. OnDemand assigns unique names to the storage objects and sequentially numbers them within an application group.

Verifying processing

The ARSLOAD program saves a copy of the messages that were generated during a load process in the system log. After a load process completes, you can open the System Log folder and review the messages. The information in the messages includes the date and time that the load process started and completed, the name of the input file, and the number of rows that were added to the database. For example:

```
arsload: Processing file >lwp<
arsload: Mon Dec 18 14:14:47 MST 1995 Indexing started,
         5116017 bytes to process
  OnDemand Load ID = >6927-0-0-4FAA-0-0<
  Loaded 4 rows into the database
arsload: Mon Dec 18 14:25:11 MST 1995 Loading Completed
arsload: Processing successful for file >lwp<
```

You can verify the number of rows that OnDemand added to the database. In the example, the number of rows added to the database should equal the number of groups (and Lotus WordPro files) that we specified in the parameter file.

The OnDemand Load ID represents the data that the ARSLOAD program stored into the system during a load process. The Load ID can be used to identify a specific load process. For example, you can run the RMVRPTOND command and specify the Load ID to delete the index data and documents that were created when the ARSLOAD program processed a Generic Indexer parameter file.

Chapter 13. Restarting a load process

The ADDRPTOND command will terminate if an unrecoverable error occurs during index, database, or storage manager processing. Termination processing includes setting a return code and saving error messages in the system log.

To start problem determination, open the System Log folder and view the messages that the command generated during the load process. The message log will contain normal processing messages, return codes, and error messages.

If the command failed during indexing, correct the problem and then restart the load process from the beginning. Common causes of problems during indexing include invalid input files or indexing parameter files and insufficient temporary space. **Note:** The messages in the system log will actually refer to a program named ARSLOAD.

If the command failed during database processing or storage manager processing:

- Determine and correct the problem.
- If a Load ID is listed in the message log that the ADDRPTOND command saved in the system log, then you can use the RMVRPTOND command to unload the data. See Chapter 14, “Deleting a report” on page 123 for information about unloading data from OnDemand.
- Restart the load process from the beginning.

The ADDRPTOND command will terminate if an unrecoverable error occurs during index, database, or storage manager processing. Termination processing includes setting a return code and saving error messages in the system log.

To start problem determination, open the System Log folder and view the messages that the command generated during the load process. The message log will contain normal processing messages, return codes, and error messages.

If the command failed during indexing, correct the problem and then restart the load process from the beginning. Common causes of problems during indexing include invalid input files or indexing parameter files and insufficient temporary space. **Note:** The messages in the system log will actually refer to a program named ARSLOAD.

If the command failed during database processing or storage manager processing:

- Determine and correct the problem.
- If a Load ID is listed in the message log that the ADDRPTOND command saved in the system log, then you can use the RMVRPTOND command to unload the data. See Chapter 14, “Deleting a report” on page 123 for information about unloading data from OnDemand.
- Restart the load process from the beginning.

© Copyright IBM Corp. 2001, 2002

Chapter 14. Deleting a report

You can use the RMVRPTOND command to delete the index data and documents that the ADDRPTOND command stored in the system during a load process. To use the delete command, you must specify the application group name and Load ID that the ADDRPTOND command generated during the load process. The Load ID represents the index data that was added to the database and the storage objects that were copied to disk and archive storage. The Load ID can be found in the message that the ADDRPTOND command saved in the System Log. If there is more than one Load ID listed in the System Log, then you should use the last Load ID listed in the log.

See online help for RMVRPTOND for more information about the command and parameters.

When the RMVRPTOND command completes, you should open the System Log folder to view the messages that were generated during the delete process. The messages will reference a program named ARSADMIN.

Note: The OnDemand server program must be running, otherwise the RMVRPTOND command will fail.

© Copyright IBM Corp. 2001, 2002

Part 5. Server administration

This section describes errors and alerts, system logging, how to find the server job and its attributes, and how to restart journaling.

© Copyright IBM Corp. 2001, 2002

Chapter 15. Managing the server

This chapter provides information for a number of topics related to the OnDemand server.

Errors and alerts

During normal processing, OnDemand programs, including the client programs, generate messages. OnDemand saves the messages in the OnDemand System Log and sends a copy of each message to the System Log user exit program. OnDemand assigns a severity to each message. Messages that are assigned a severity of alert or error are also sent to the OS/400 QSYSOPR message queue and logged in the server joblog. Additionally, when a user runs a query that requires a table of index data that has been migrated to archive storage, OnDemand sends a message to QSYSOPR and the joblog.

OnDemand provides the System Logging facility and a message catalog to help you identify and resolve any alerts and errors that you may receive. You can open the System Log folder to display the messages that are saved in the System Log. See "System logging facility" for more information.

System logging facility

OnDemand provides a logging facility to help administrators track OnDemand activity and monitor the system. When you enable logging for system events, user events, and application group events, OnDemand stores the messages that are generated by the various OnDemand programs in the System Log. You can use the OnDemand client program to search for and filter messages by time stamp, severity, message number, and user name.

Searching for and viewing messages

To search for and view the messages that are stored in the System Logging facility, log on to OnDemand with an OnDemand client program and open the System Log folder. Enter search criteria in one or more of the search fields. When you choose the **Search** command, OnDemand retrieves the messages from the database that match the search criteria that you specified.

You can specify a date and time value to search for and retrieve messages. You can also specify other search criteria, such as:

Log Id Each time that a client logs on to the server, OnDemand assigns a number to that session. All messages that are generated during that session include the same Log Id.

Userid The OnDemand userid

Severity OnDemand assigns a severity to each message: Alert, Error, Warning, Info, and Debug

View Depending on the type of message in the System Log, you may be able to view other information that is related to or associated with the message. For example:

- You can display the message log that was generated during a load process by selecting an ARSLOAD message and then

© Copyright IBM Corp. 2001, 2002

choosing the **View all Selected** command. You might see ARSLOAD messages after issuing the ADDRPTOND command.

While most processes do not generate other information that can be stored in the System Log, you could write a user exit program to process the messages and generate your own information about the events. For example, you could write a user exit program to generate a report that lists the number of users that are logged on to the system in 30 minute increments. OnDemand provides a System Log user exit so that you can process any message that is stored in the System Log and take the action that you require. See below for more information about the System Log user exit.

- Other messages in the System Log do not provide additional records or other data associated with an event. For example, the Logon and Logoff events each generate a single message, with no additional information that you can view.

Msg Num The message number that is assigned by OnDemand

Message The text of the message that OnDemand uses to restrict a search. For example, if you type **Login**, OnDemand searches for and displays the messages issued by the Logon to a Server command.

System Log user exit

When you enable logging for system, user, and application group events, OnDemand sends a copy of each message that is generated by the system to the System Log user exit program. The System Log user exit program is named `arslog` and resides in the OnDemand library named QRDARS.

The System Log user exit program that is provided by IBM does not perform any functions. However, you can replace the program that is provided by IBM with your own program that does user-defined processing. For example, you could create a program that checks for certain message numbers or severity and takes whatever action you deem appropriate.

You configure OnDemand to send messages to the System Log user exit by selecting User Exit Logging options with the System Parameters command. See the online help for the administrative client for more information about the User Exit Logging options and the System Parameters command.

See the *IBM Content Manager OnDemand for iSeries Common Server Planning and Installation Guide* for more information about the System Log user exit.

Monitoring users

The OnDemand server generates System Log messages to help you track the number of users that are logged on to the server:

- OnDemand stores message number 201 in the System Log approximately every 30 minutes if at least one activity is detected. This message contains the current number of users that are logged on to the server.
- OnDemand stores message number 202 in the System Log every time that the number of concurrent users exceeds the previous maximum number of concurrent users. The number of concurrent users is reset each time that you restart the OnDemand server processes.

Reviewing a CD-ROM creation manifest

You can open the System Log folder to review a manifest created by the OnDemand ad-hoc CD-ROM mastering software. A CD-ROM creation manifest contains a cover page with control information about the CD-ROM and one or more pages that list the items retrieved from the database and written on the CD-ROM.

To review a CD-ROM creation manifest:

1. Log on to OnDemand with an OnDemand client program and open the System Log folder.
2. Type **CD** in the Message field.

   Note: If you know the CD-ROM volume number of the manifest that you want to review, then enter the number in the Message field. For example, you can enter **%171** in the Message field to locate the manifest for volume number AOD00000171.
3. Specify a date or date range if desired.
4. Choose **Search**. OnDemand displays the manifests that match the search criteria that you specified.
5. Review a manifest by selecting an item from the list and choosing the **View All Selected** command.

How to find the server job and its attributes

The name of the server job matches the name of the OnDemand instance for which it is running. For example, the server job name is QUSROND for the the default instance that is created when OnDemand is installed. You can use the command WRKACTJOB JOB(QUSROND) to find this default server job.

Server jobs are started using a job description by the name of the instance (which must be found in the QUSRRDARS library). If a job description by that name is not found in QUSRRDARS, then job description QOND400 in library QRDARS is used (and can be changed if necessary).

The job description controls the following attributes of the server job:

- JOBQ
- JOBPTY
- OUTPTY
- PRTDEV
- OUTQ
- INLLIBL
- LOG
- LOGCLPGM
- INQMSGRPY
- HOLD
- DATE
- SWS
- JOBMSGQMX
- JOBMSGQFL

For example, if you wanted to change the job queue that instance TEST used, you would create a job description called TEST in library QUSRRDARS that specified the job queue you wanted to use.

Restarting journaling

If you receive a message on the server indicating that journaling needs to be started for your OnDemand database files, issue the following command from a command line while signed on with sufficient authority:

CALL QRLCSTRJ PARM(RLC)

The parameter value (RLC) must be entered in *uppercase*.

Part 6. Appendixes

© Copyright IBM Corp. 2001, 2002

Appendix A. Command reference

This part of the book contains reference information about the OnDemand server commands. The commands are presented in alphabetical order. Each command contains a description of its purpose. See the online help text for details about the command parameters.

OnDemand server commands

These commands require that you be signed on to the server with a user profile that is also defined as a user in OnDemand. For commands that specify an application or application group name, if the value to be entered contains lower case letters, blanks, or special characters, it must be enclosed in apostrophes.

ADDRPTOND

The Add Report to OnDemand (ADDRPTOND) command allows you to load reports in OnDemand. During this process the report is broken into segments, indexed, compressed, and stored on disk for retrieval and later migration to optical or tape media if desired.

ENDMONOND

The End Monitor for OnDemand (ENDMONOND) command allows you to end a currently active monitor for an output queue.

FNDKEYOND

The Find Key for OnDemand (FNDKEYOND) command allows you to search for a particular document that is available through the folder specified in the FOLDER parameter. The OnDemand client is launched to view the results of the search. The key fields entered must exist in the folder. This is intended as an API to launch the OnDemand client from a green screen application.

PRTRPTOND

The Print Report from OnDemand (PRTRPTOND) command prints the specified report in its entirety.

If a printer name is specified for the PRINTER parameter, then OnDemand will use the first record found in the ARSPRT file (in your instance library) with a printer name that contains the string you entered anywhere in the name.

For example, if you enter PRINTER1 for the PRINTER parameter and your ARSPRT file has these four records (in this order), OnDemand would send your output to the first one in the following list:

- Printer1 in the printer room
- PRINTER1
- Printer123
- XPrinter1

If *OUTQ is specified for the PRINTER parameter, then you must also enter an output queue name and library. OnDemand will use the printer associated with the first record found in the ARSPRT file (in your instance library) with an output

© Copyright IBM Corp. 2001, 2002

queue name that contains the library name/output queue name you entered. If no match is found and if *LIBL is entered as your output queue library name, OnDemand will use the printer associated with the first record found in the ARSPRT file (in your instance library) with an output queue name that matches (with no library specified).

For example, if you enter QUSROND/MYOUTQ for the OUTQ parameter and your ARSPRT file has these four records (in this order), OnDemand would send your output to the printer name associated with the first output queue in the following list:

- QUSROND/MYOUTQ
- MYOUTQ
- AMYOUTQ
- TEST/MYOUTQ

PRTTXTOND

The Print Text for OnDemand (PRTTXTOND) command allows you to print all or part of a spooled file in a text-only format. This allows the report administrator to see what the report looks like to ADDRPTOND when it is indexed. Using this output, the administrator can determine how to index the report.

RMVRPTOND

The Remove Report from OnDemand (RMVRPTOND) command removes the specified report from OnDemand.

STRASMOND

The Start Archived Storage Mgmt (STRASMOND) command starts the Archived Storage Management (ASM) process which manages the movement of data within the ASM defined levels. This function can be canceled with a controlled cancel option if enough time is allowed to end what it is currently processing.

Notes:

1. This command must only be run in batch (SBMJOB parameter set to *YES). Running this command interactively (with SBMJOB(*NO)) may cause SQL errors.
2. The STRASMOND command does not have a specific parameter for specifying an instance name. By default, the QUSROND default instance is used, and will produce the desired results for most systems. If you need to name a different instance, you must specify it as INSTANCENAME/POLICYNAME in the Policy parameter of the STRASMOND command. You do not need to put your entry in apostrophes. If you enter an instancename/policyname, the command prompter will automatically enclose it in apostrophes. This will force the value to be processed in whatever case it was entered (uppercase, lowercase, or mixed case). Note that if you initiate the archive storage manager by running the STRDSMOND command with RUNASM(*YES), then the instance name is passed from the disk storage manager and no further specifications are needed.

STRDSMOND

The Start Disk Storage Management (STRDSMOND) command starts the Disk Storage Management (DSM) task which manages the movement of OnDemand data on disk and between disk and the Archived Storage Manager (ASM). DSM

also controls the expiration of data in OnDemand. This function can be canceled with a controlled cancel option if enough time is allowed to end what it is currently Processing.

Note: If you run STRDSMOND for a specific application group (rather than the default of *ALL) and you set the Run ASM (RUNASM) parameter to *YES, be aware that ASM will run for ALL application groups, even though you have named a specific application group for DSM to use. You can, however, name a specific Policy for ASM to process, if desired. Also note that when you specify RUNASM(*YES), OnDemand will initiate a separate batch job for ASM.

STRIMPOND

The Start Import into OnDemand (STRIMPOND) command allows you to import data into OnDemand. This command is used only if you have migrated your index data to an alternate media (such as optical or tape), which is not recommended, but may be necessary in some cases.

STRMONOND

The Start Monitor for OnDemand (STRMONOND) command allows you to specify the name of an output queue or IFS directory to monitor. When a spooled file is added to the output queue or a file is added to an IFS directory, the spooled file is automatically processed by ADDRPTOND.

The application and application group names for spooled files are determined using spooled file attributes such as spooled file name or user data (or a number of others). You can alter these values that OnDemand uses for application and application group by using the output queue monitor user exit program as described in Appendix D, "API and user exit reference" on page 145.

The application and application group names for PC files are determined by the PC file name itself. The monitor can use any one of the first four parts of the file name. For example, a file name of `CHECKSTMTS.BIGSTMTS.data` might be used to archive an application named `CHECKSTMTS` into an application group named `BIGSTMTS`.

You can end the monitor by:

- Specifying a method for the monitor to end automatically.
- Running the ENDMONOND command
- Ending the monitor job using the ENDJOB command. Specify OPTION(*CNTRLD) and DELAY(999999). The job will end as soon as the monitor finishes processing the current file.

Appendix B. 5250 host connection to client viewer

This section describes the OnDemand 5250 Host Connection, which allows an OS/400 application to send information to the OnDemand client workstation viewer program (the viewer). The intent is to provide a mechanism by which a 5250 line-of-business application can use the viewer to display documents from the OnDemand database. This is done with little or no interaction between the user and the viewer.

This function is comprised of two parts:

- The FNDKEYOND or FNDKEYRDAR command, which the line-of-business application will invoke. The FNDKEYOND command is a part of Option 10 (Common Server) of OnDemand for iSeries, and is used to retrieve documents archived using the OnDemand Common Server feature. The output of the FNDKEYOND command MUST be displayed using the workstation viewer. FNDKEYRDAR is a part of Option 1 (Spool File Archive) of OnDemand for iSeries, and it is used to retrieve documents archived using the OnDemand Spool File Archive feature. The output of the FNDKEYRDAR command may be displayed using the workstation viewer, or may be displayed on a non-programmable terminal or in a terminal emulation session. The references to FNDKEYRDAR which follow assume that its output is being directed to the workstation viewer. One or both of these commands may be used, depending on which options of OnDemand (Spool File Archive or Common Server) are installed. In the following paragraphs, FNDKEYxxxx means either FNDKEYOND or FNDKEYRDAR.
- A companion program that resides on the workstation, which is the interface between the FNDKEYxxxx command and the viewer.

Operational and environmental considerations

For this function to operate correctly, you must address several operational and environmental considerations. This section will describe these considerations.

Workstation installation tasks

For the FNDKEYxxxx command to operate correctly, the workstation companion program, QRLROCD.EXE, must be running on the workstation when the FNDKEYxxxx command is run on the server. This program is found in the OnDemand install directory, usually C:\Program Files\IBM\OnDemand32. The following paragraphs describe a suggested approach for simplifying the startup of this program.

After the OnDemand client software is installed on your workstation, you should establish a shortcut which points to the QRLROCD.EXE program. This shortcut can be placed on the workstation desktop, in the startup folder or in any other place which meets your needs. Placing the shortcut in your startup folder will start the program automatically when you start your workstation. This is appropriate if you use this interface frequently. If you end the program and need to restart it, or if you have no need for the program to be started automatically, you may want to place a shortcut on the desktop.

In most cases, you will only need to start the companion program once. If, however, the workstation is attached to multiple OnDemand server systems, or if

© Copyright IBM Corp. 2001, 2002

both FNDKEYOND and FNDKEYRDAR output will be displayed by the viewer, you may need to start the program multiple times. In this case, you should create a shortcut for each one, for reasons explained below in the section labeled Multiple OnDemand servers in the network. Follow the directions in the sections below, depending on your configuration. The term ″multiple OnDemand server systems″ refers to a network where either of the following is true:

- there are two or more OnDemand server systems from which you need to view FNDKEYxxxx output concurrently using the viewer
- there is one OnDemand server, but the results of both FNDKEYOND and FNDKEYRDAR commands will be displayed concurrently by the viewer

One OnDemand server in the network

If you will be using either FNDKEYOND or FNDKEYRDAR (but not both), and there is only one OnDemand server where the command will be used, you may allow the command and the companion program to use the default port number (3005). In this case, it is not necessary to modify the shortcut to specify a different port. Please skip to the section titled Other workstation considerations. If the above does not describe your configuration, please continue with the next section.

Multiple OnDemand servers in the network

If you intend to display the results of both the FNDKEYOND and FNDKEYRDAR commands with the viewer, or if there are multiple OnDemand servers where either of these commands may be run, a separate workstation companion program must be started on the workstation for each different combination of system and command. For example, if FNDKEYRDAR and FNDKEYOND will be run on a single system, two of the companion programs should be started. The same applies when FNDKEYOND is run from two systems.

Each instance of the companion program must be configured to use a different port. After you have created the shortcuts you need, you should alter the properties of the shortcuts to specify a port number and, optionally, a system name. The system name, if specified, will appear as the window title for the viewer which the shortcut will start.

To specify a port and system for a shortcut, right-click the shortcut. From the pop-up menu, select Properties. On the Shortcut tab, you will see an entry box labeled Target. This should contain the path for the QRLROCD.EXE program. At the end of the path, after QRLROCD.EXE, add at least one space, followed by:

```
/p=nnnn  /s=systemName
```

where nnnn is the port number and systemName is the name of the OnDemand server.

Example:

```
..../qrlrocd.exe   /p=3006  /s=ACCOUNTING
```

You may also use uppercase P and S (...qrlrocd.exe /P=3007 /S=BILLING).

After adding the port and system parameters, press OK to save the new properties of the shortcut. You also should change the title which displays beneath the shortcut to something meaningful, such as ″OnDemand Viewer for ACCOUNTING reports″.

Make a note of the port numbers and system names you have specified for the shortcuts. It will be necessary to configure the server systems so that FNDKEYxxxx will use the matching port numbers instead of the default. This is discussed in the section ″Server Configuration″ below.

Other workstation considerations

For FNDKEYxxxx output to be displayed correctly using the OnDemand viewer, the FNDKEYxxxx program on the server system must be able to determine the TCP/IP address of the workstation. You should access your line-of-business applications which invoke FNDKEYxxxx through a terminal emulation session. The connection method for the session should be TCP/IP. If the FNDKEYxxxx command is run in a job attached to this session, the address of the workstation can be determined by the FNDKEYxxxx command, and nothing further needs to be done.

If the emulation session does not use TCP/IP, the FNDKEYxxxx command searches for the presence of a data area, located in QUSRRDARS library, which has the same name as the device name associated with the emulation session. If the data area exists, the command looks for a non-blank value in positions 1-15 of the data area and uses this value as the workstation IP (Internet Protocol) address. It is the responsibility of the user to ensure that this data area exists and contains the IP address. You can use the CRTDTAARA and CHGDTAARA commands for this purpose.

If the FNDKEYOND command detects errors when sending the search request to the workstation companion program, it will assign a three-digit error code and will display this information in message RDR2882. If the FNDKEYRDAR command detects errors, see the *IBM Content Manager OnDemand for iSeries Common Server Administration Guide* for an explanation of the errors.

Table 6 lists the error codes and their meanings for exception RDR2882.

Table 6. Exception RDR2882 error codes

Error code	Error description
002	Environment error. A connection could not be established with the workstation companion program because the workstation IP address could not be determined.
091	The TCP/IP socket could not be created.
092	Connect failed for TCP/IP socket. The most likely cause is that the workstation companion program is not running, or there is a port number mismatch between the FNDKEYxxxx command and the companion program.
093	Socket Close operation failed.

Server Configuration

The information in this section is pertinent only if you created more than one shortcut. If you created only one shortcut, the FNDKEYxxxx commands will function correctly using the default port number, and nothing further needs to be done.

If you created more than one shortcut, you should also have specified different port numbers for the shortcuts. Since the FNDKEYxxxx commands must use the same port numbers which you specified on the shortcuts, it is necessary to configure these commands on the server systems to use the specified ports.

The FNDKEYxxxx commands check for the existence of a data area named QRLMFKPORT, in library QUSRRDARS. If this data area does not exist, the commands will use the default port number (3005). If the data area exists, the FNDKEYxxxx commands will use the port number found in the data area contents. The first five characters of the data area contents represent the port number which the FNDKEYOND command should use. The next five characters represent the port number for the FNDKEYRDAR command. To create this data area, enter the following:

```
CRTDTAARA  QUSRRDARS/QRLMFKPORT TYPE(*CHAR)  LEN(10)  VALUE( 'AAAAABBBBB' )
```

where AAAAA is the FNDKEYOND port number and BBBBB is the FNDKEYRDAR port number.

For example, to specify that FNDKEYOND, on this system, should use port 3007 and that FNDKEYRDAR should use 3008, specify VALUE('0300703008'). It is important that each five-position number be right-justified, with leading zeros if necessary. If you will not use both commands from the same system, enter 00000 for the port number corresponding to the command you will not use.

If there are multiple OnDemand servers in your network (or you will use FNDKEYOND and FNDKEYRDAR commands on one or more systems), and there are workstations which will display FNDKEYxxxx output from these systems, you must configure the FNDKEYxxxx commands on these systems to use a different port number. Some examples are shown below:

- If you have one OnDemand server, but both FNDKEYOND and FNDKEYRDAR commands will be used on this system, you need to create the data area. In this case, you can allow one of the commands to use the default port, and the other command should use a different port. For example, if you specify VALUE('0300503006'), FNDKEYOND will use port 3005 , and FNDKEYRDAR will use 3006. Make sure that your shortcuts are modified to specify these port numbers.
- If FNDKEYxxxx will be run from multiple servers in your network, create the data area on the systems where it is necessary to change the port assignments for the FNDKEYxxxx commands. For example, on SYSTEMA, specify VALUE('0300503006'), and on SYSTEMB specify VALUE('0300703008'). It is imperative that these commands on multiple systems in the network use different port numbers.

Appendix C. Server printing and faxing

Using an OS/400 printer file to define server print parameters

You can use an OS/400 printer file to define some of the server print parameters in your Common Server application definition by specifying the following in the Print Parameters field on the Print Options tab of your application:

```
PRTF=LIBRARY/PRINTERFILENAME
```

Where `LIBRARY` is the name of the library that contains your printer file and `PRINTERFILENAME` is the name of the printer file.

Note: The entire line must be entered in uppercase.

This printer file specification is also used when faxing from the server to produce the temporary OS/400 spooled file which is then faxed.

The following parameters, even though specified in the printer file identified in the PRTF parameter, are overridden as shown in Table 7.

Table 7. Parameter Overrides

Printer file parameter	Changed to	Where specified	When specified
DEVTYPE	*AFPDS if AFP data. *SCS if SCS or SCS-extended data. *LINE if Line data and an EBCDIC code page is used; otherwise it is set to *SCS and the data is converted from ASCII to EBCDIC.	Application definition - View Information tab - Data Type field	When defining the application
OUTQ	Server Queue Name	Printer definition	When printing the document, you select the Server Printer definition to use
COPIES	Number of Copies	Client Print window	When printing the document
USRDTA	The first 10 characters of your OnDemand userid	When you logon to OnDemand	When you logon to OnDemand
USRDFNDTA	Application name	Application definition - General tab - Name field	When defining or updating the application
PAGERANGE	Pages value	Client Print window, but only if already viewing the document	When printing the document

If no PRTF parameter is specified, printer file QSYSPRT is used and in addition to the printer file parameters specified above, the parameters in Table 8 on page 142

© Copyright IBM Corp. 2001, 2002

are also set.

Table 8. QSYSPRT Parameters

Printer file parameter	Changed to	Where specified	When specified
CTLCHAR	*FCFC for Line data; otherwise use QSYSPRT (default) **Note:** If the View Information tab in the Application definition specifies CC - No, the data is printed using an ANSI carriage control character of blank for each line (single spaced).	Application definition - View Information tab - Data Type field	When defining or updating the application
PAGRTT	Orientation (for AFP)	Application definition - View Information - Orientation field	When defining or updating the application
TBLREFCHR	TRC (for Line data)	Application definition - View Information - TRC field	When defining or updating the application

Server fax setup

When defining a printer for use with server fax functions, a Server Queue Name is required in the printer definition within the OnDemand Administrator. You can either specify the name of a library/output queue or *NONE. If *NONE is specified, the output is temporarily spooled to QUSRRDARS/QRDARS400 output queue before it is sent to fax using the QRLMSFAX program. If you enter a library/output queue name, then OnDemand will use that output queue to temporarily spool the data.

If you use the SNDFAX command in the QRLMSFAX program (which is the default as shipped), the restrictions in Table 9 apply to the fax information. If you use a different server fax product, see "Facsimile user exit program" on page 178.

Table 9. SNDFAX command fax information restrictions

Administrator Client Prompt Field	SNDFAX maximum length	SNDFAX keywords used	Length passed to QRLMSFAX
Recipient Attention	40	TO position 2	100
Recipient Company	40	TO position 3	100
Recipient Fax Number	32	TO position 1, TO position 4	100
Sender Name	40	FROM position 1	100
Sender Company	40	FROM position 2	100
Sender Tel Number	40	FROM position 3	100
Sender Fax Number	–	(not used)	100

Table 9. SNDFAX command fax information restrictions (continued)

Administrator Client Prompt Field	SNDFAX maximum length	SNDFAX keywords used	Length passed to QRLMSFAX
Sender Cover Page	10	CRTCVRP - can be blank, *NO or *YES; default is *YES	10
Subject	40	TITLE	100
Notes	40	COMMENT	100

Appendix D. API and user exit reference

API reference

This section contains reference information about the OnDemand server Application Programming Interfaces (APIs).

The APIs contained in this reference include APIs for various OnDemand functions beyond the standard command set described in Appendix A, "Command reference" on page 133. The APIs are presented in alphabetical order. Each API contains a description of its purpose and syntax (including descriptions of the parameters that can be used). Examples and general information about using the APIs have also been included.

Note: These APIs require that you be signed on to the server with a user profile that is also defined as a user in OnDemand.

Using quotes when executing the APIs

When you execute an OnDemand program from the OS/400 command line interactively, in batch using the SBMJOB command or the QSHELL environment and you specify parameter values that contain a null (blank) character or some other special character (such as the parenthesis), you must delimit the parameter value with double quote characters. For example, when executing the ARSDOC GET program from the OS/400 command line and you specify the **-o** parameter to specify one or more application group field names, each field name is enclosed with parenthesis and the entire string must be surrounded by double quote characters, for example: `-o "(sdate)(student)"`. Please see the operating system documentation for more information about the use of quotes when executing the APIs. Note that if the same string is included in a parameter file using the **-F** parameter on the ARSDOC GET command, it would be enclosed in brackets and the double quotes are not required, for example: `[-o (sdate)(student)]`.

Using the QSHELL environment

All **ARSxxxxxx** API programs must be executed using the QSHELL environment. QSHELL is a command interpreter that allows OS/400 to execute AIX® commands on OS/400. These commands can also be run from a script file. More information about QSHELL can be found at `http://www.as400.ibm.com/developer/qshell/`. Currently, option 30 for 5722-SS1 (OS/400) installs the QSHELL intepreter. The QSHELL environment is started by the **QSH** or **STRQSH** OS/400 commands.

Multiple **ARSxxxxxx** programs may be issued in one QSHELL session. After the completion of a QSHELL command, a $ is displayed. This notifies the user that the QSHELL command has completed. This does not mean that the command actually ran successfully and performed the desired function.

All command examples assume that the QSHELL environment has been started and the **ARSxxxxx** programs are in a directory specified in the PATH variable. OnDemand places a symbolic link to the commands in the `/usr/bin` directory during installation. To exit the QSHELL environment, press the F3 key.

The **-v** parameter will provide detailed program logging information. The **-u** parameter specifies a valid OnDemand user that exists with the proper authority to

© Copyright IBM Corp. 2001, 2002

perform the functions requested. The **-p** parameter specifies the OnDemand password for the OnDemand user specified in the -u parameter. The userid and password are not normally needed when running these programs in OS/400. By default, the current OS/400 user profile and password are used as the OnDemand userid. If your OS/400 user profile does not exist in OnDemand, you must specify a valid OnDemand userid and password to use these programs.

The **-h** is the OnDemand instance name where the program is to execute. For the purposes of this document, the administrative user name will be `testadmin`, the password will be `ondemand` and the host name is `QUSROND`. User `testadmin` must have the appropriate authority to perform the function being requested by the **ARSxxxxxx** command.

Calling QSHELL commands from an OS/400 command line

When executing the **ARSxxxxxx** programs from an OS/400 command line, the OS/400 QSH command is used. When embedded single quotes exist within the program parameter, they must be doubled. For example:.

```
QSH CMD('arsdoc get ..... "-i WHERE Account#=''1234567'' "....-u testadmin
-p ondemand')
```

Since the `QSHELL` environment will log messages to the terminal, you will want to control this by setting environment variable `QIBM_QSH_CMD_OUTPUT`. This variable will control where the messages are logged. This does not interfere with the messages that get logged to the OnDemand System Log. The QSHELL environment logging can be sent to the terminal session, to a file in the IFS directory structure, or you can choose not to log the messages. The ADDENVVAR, CHGENVVAR, RMVENVVAR and WRKENVVAR commands can be used to manipulate the `QSHELL` environment. The environment parameters can be set for the entire system and for the job. In most cases, you will not want to change the environment parameter at the system level, because it will affect all users of `QSHELL` on the system.

- Set the job environment to log messages to the terminal
  ```
  ADDENVVAR ENVVAR(QIBM_QSH_CMD_OUTPUT)  VALUE(STDOUT) LEVEL(*JOB)
  ```
- Set the system environment to not log messages.
  ```
  ADDENVVAR ENVVAR(QIBM_QSH_CMD_OUTPUT) VALUE(NONE) LEVEL(*SYS)
  ```
- Set the job environment to log messages to the a file in the IFS. mydirectory must exist in the IFS.
  ```
  ADDENVVAR
  ENVVAR(QIBM_QSH_CMD_OUTPUT)VALUE('FILE=/mydirectory/QSHELL_output')
  LEVEL(*JOB)
  ```

The following is an example of the messages that might be displayed to the terminal or logged in the output file.

```
OnDemand Load Id = >5013-1-0-85FAA-11359-11359<
Loaded 4 rows into the database
Document compression type used - OD77.  Bytes Stored = >10240<
```

Message QSH0005 will be issued in the joblog when running the QSH command. The message text is ″Command ended normally with exit status &1″. The possible statuses returned are shown in Table 10 on page 147:

Table 10. Message QSH0005 status codes

Status	Description
0	Completed Successfully
1	Command Failure
2	Folder does not exist or do not have authority to folder
3	User ID or Password is not valid. Cannot establish communication to server
127	Command Not Found
254	No Hits Match Query

How to read a syntax diagram

A syntax diagram shows you how to specify an API program so that the operating system can correctly interpret what you type.

Read a syntax diagram from left to right and from top to bottom, following the horizontal line (the main path). If the line ends with an arrowhead, the API syntax is continued and the next line starts with an arrowhead. Facing arrowheads mark the end of the API syntax.

When you type an API from the syntax, be sure to include punctuation, such as commas and equal signs.

Parameters are classified as keywords or variables:

- Keywords represent constants and are shown (in syntax) in uppercase letters; however, you can enter keywords in either uppercase or lowercase.
- Variables represent names or values you supply and are shown (in syntax) in lowercase letters; however, you can enter variables in either uppercase or lowercase unless the syntax directions explicitly state the case restrictions. An example of a variable is a file name.

A parameter can be a combination of a keyword and a variable.

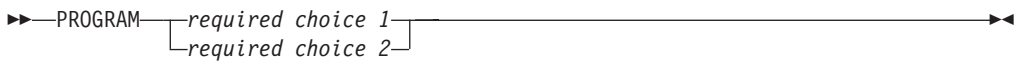
Required parameters are displayed on the main path.



Optional parameters are displayed below the main path.



A stack of parameters, with the first parameter displayed on the main path, shows that you must choose one of the parameters.



A stack of parameters, with the first parameter displayed below the main path, shows that you can choose one of the parameters.

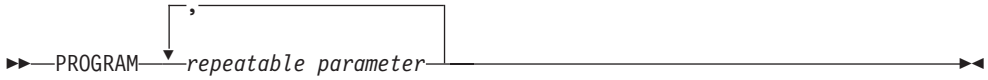


An arrow returning to the left, above the path, shows that items can be repeated following these conventions:

- If the repeat arrow contains a break, the item can be repeated in a list with the items separated by blank spaces.

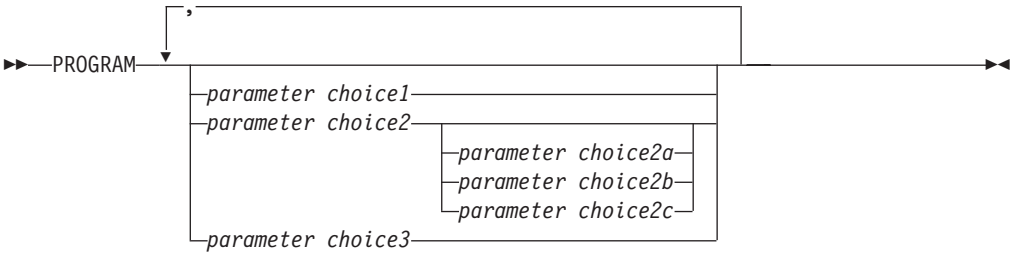


- If the repeat arrow contains a comma, the item can be repeated in a list with the items separated by commas.



You can repeat items from parameter stacks following the stack conventions for required and optional parameters described previously.

Some syntax diagrams contain parameter stacks within other parameter stacks. You can only repeat items from stacks according to the conventions described previously. That is, if an inner stack does not have a repeat arrow above it but an outer stack does, you can choose only one parameter from the inner stack and combine it with any parameter from the outer stack, and that combination can be repeated. For example, the following diagram shows that you could combine parameter choice2a with parameter choice2 and then you can repeat that combination again (choice2 plus choice2a).



Some APIs are preceded by an optional path parameter.



If you do not supply the path parameter, the system searches the current directory for the API. If the API is not in the current directory, the system continues to search for the API using the directories defined in the PATH environment variable.

Some APIs in this section have several formats that accomplish the same task. These APIs appear (in syntax) similar to the following:



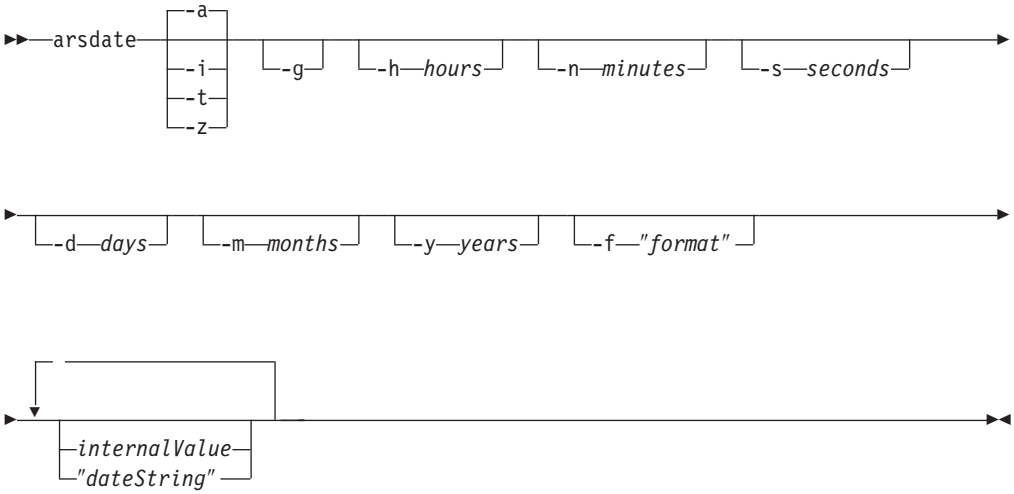
The description of the API directs you to the correct format to use.

ARSDATE

Purpose

The ARSDATE program allows you to display the OnDemand internal database value for a given date and time string or display the date and time string for a given OnDemand internal database value.

Syntax



Description

The ARSDATE program can be used to display the OnDemand internal database value for a given date and time string or display the date and time string for a given OnDemand internal database value.

Values whose data types are Date, Time, Date/Time, or Date/Time (TZ) are represented in an internal form that is transparent to the casual user of OnDemand. Casual users enter these values the same way that they appear in a report. However, to search the database with an SQL string, a user must enter the internal form of the value. The ARSDATE program lists the internal value of a date or time string.

Parameters

-a Use to display database values and date strings for Date fields. For example, to display the database value for the date 9/1/99, enter:

```
arsdate -a 9/1/99
```

The ARSDATE program returns:

```
9/1/99 -> 10836
```

To display the date string for the database value 10836 in the default date format, enter:

```
arsdate -a 10836
```

The ARSDATE program returns:

```
10836 -> 9/1/99
```

-i Use to display database values and date and time strings for Date/Time fields. The time part of a Date/Time field is not adjusted for the local time zone. You typically use the **-i** parameter to find out one of two things:

- Given a date and time printed on a report, what value did OnDemand store in a Date/Time database field? You can use the result to search a Date/Time field with an SQL string.
- Given a value stored in a Date/Time database field, what would be the date and time printed on a report?

-t Use to display database values and time strings for Time fields. The time is not adjusted for the local time zone. For example, to display the database value for the time 04:00:00, enter:

```
arsdate -t 04:00:00
```

The ARSDATE program returns:

```
04:00:00 -> 4800
```

To display the time string for the database value 4800, enter:

```
arsdate -t 4800
```

The ARSDATE program returns:

```
4800 -> 04:00:00
```

-z Use to display database values and date and time strings for Date/Time (TZ) fields. The time part of a Date/Time (TZ) field is adjusted for the local time zone. If you run the ARSDATE program with the **-z** parameter on systems in different time zones and you specify the same date and time value, the result will be different. For example, suppose that you need to determine the value stored in a Date/Time (TZ) field for ″09/01/99 04:00:00″. The API:

```
arsdate -z "09/01/99 04:00:00"
```

When run on a server in the Eastern time zone will return:

```
09/01/00 04:00:99 -> 936172800
```

If you run the same API on a server in the Mountain time zone, then result will be:

```
09/01/00 04:00:99 -> 936180000
```

A typical use of the **-z** parameter is to determine a database value with which to search the system log. You can use the result to search the Date/Time (TZ) field of the system log with an SQL string. For example, suppose a user in New York logs on to a server in Denver. To retrieve the log on messages with an SQL string, you must specify the date and time part of the query using the local time of the client that is running the query. If the user

logged on to the server at 4 a.m. Eastern time, then a query that is run in Denver must specify 2 a.m. to retrieve the message.

-g Display the current system date.

-d days Add the specified number of days to the specified internal value.

-h hours Add the specified number of hours to the specified internal value.

-m months Add the specified number of months to the specified internal value.

-n minutes Add the specified number of minutes to the specified internal value.

-s seconds Add the specified number of seconds to the specified internal value.

-y years Add the specified number of years to the specified internal value.

-f "format" Determines the format of the date and time string that OnDemand displays.

The following table lists some of the formats supported by OnDemand. If the format that you require is not listed, you can create your own, using values, separators, and the percent character.

The format consists of a set of values (for example, **%m**) and separators (such as the blank character).

Table 11. Date and time format specifications

Date Format Specifier	Date Format	Example
%m/%d/%y	mm/dd/yy	01/31/95
%d/%m/%y	dd/mm/yy	31/01/95
%f/%e/%y	m/d/yy	1/31/95
%e/%f/%y	d/m/yy	31/1/95
%m-%d-%y	mm-dd-yy	01-31-95
%d-%m-%y	dd-mm-yy	31-01-95
%m%d%y	mmddyy	013195
%m%d%Y	mmddyyyy	01311995
%y%m%d	yymmdd	950131
%Y%m%d	yyyymmdd	19950131
%b %e, %Y	Mth d, yyyy	Jan 31, 1995
%B %e, %Y	Month d, yyyy	January 31, 1995
%m/%d/%y %H:%M	mm/dd/yy hh:mm	01/31/95 10:50
%H:%M	hh:mm	10:50
%T	hh:mm:ss	10:50:59

internalValue The internal date and time value from the OnDemand database. Enter one or more internal values when you want to display formatted date strings.

"**dateString**" The date and time string. Enter one or more strings when you want to display internal date values.

Examples

1. The following example shows how to determine the database value for the specified date and time string. The data type of the database field is Date/Time.

```
arsdate -i "09/01/99 04:00:00"

09/01/99 04:00:00 -> 936158400
```

   In the example, you could use the database value to search a Date/Time field in the database with an SQL string. For example:

```
arsdoc get -i "WHERE somedate=936158400" ...
```

2. The following example shows how to determine the date and time string for the specified database value. The data type of the database field is Date/Time. The result is shown using the default display format.

```
arsdate -i 936158400

936158400 -> 09/01/99 04:00:00
```

3. The following example shows how to determine the database value for the specified date and time string. The data type of the database field is Date/Time (TZ). The ARSDATE program adjusts the time part of the result for the local time zone. For Mountain time zone:

```
arsdate -z "09/01/99 04:00:00"

09/01/99 04:00:00 -> 936180000
```

   If you were to run the same API on a server in the Eastern time zone, the result would be:

```
arsdate -z "09/01/99 04:00:00"

09/01/99 04:00:00 -> 936172800
```

Notes

1. When displaying the date string for a given internal value, by default, the ARSDATE program displays the date string using the `mm/dd/yy` format. If you want the ARSDATE program to display the date string using a different format, then you must specify the date format with the **-f** parameter. For example:

```
arsdate 10907

10907 -> 11/11/99

arsdate -f "%m/%d/%Y" 10907

10907 -> 11/11/1999
```

2. When displaying the internal value for a given date string, by default, the ARSDATE program expects you to specify the date string using the `mm/dd/yy` format. If you want to specify the date string using a different format, then you must specify the date format with the **-f** parameter. For example:

```
arsdate 11/12/99

11/12/99 -> 10908

arsdate 11/12/1999

11/12/1999 -> -1 (Error)

arsdate -f "%m/%d/%Y" 11/12/1999

11/12/1999 -> 10908
```

IFS location

/usr/bin/arsdate The executable program.

ARSDOC

Purpose

The ARSDOC program is a multi-purpose document processing program. You can use the ARSDOC program to query the server and generate a list of items that match a query, retrieve documents from the system, add, delete, and update documents, and send documents to a server printer. You can run the ARSDOC program from the command line or from a program. The ARSDOC program provides the following functions:

Description and syntax

The ARSDOC program provides the following functions:

- ADD
- DELETE
- GET
- PRINT
- QUERY
- UPDATE

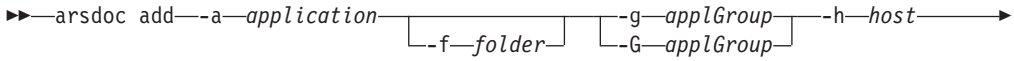
For each function, you can specify all of the required options on the command line or you can specify the name of a parameter file that contains the options. We've listed the syntax of each function twice: first, when you do not use the parameter file option, then the syntax when you specify a parameter file (along with the syntax of the parameter file).

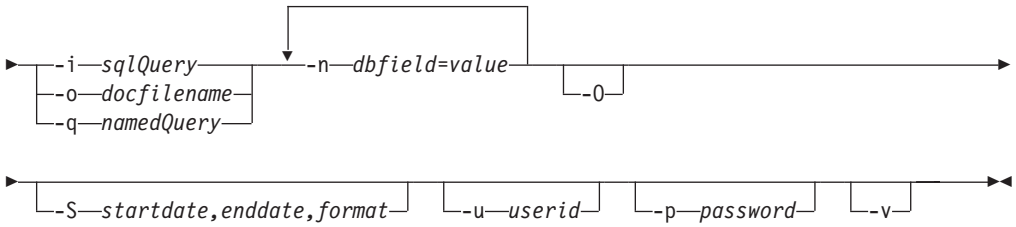
ADD function: Use to store data in OnDemand by specifying the folder, application group, application, and database fields and values. If you specify the **-O** parameter, then you do not have to specify all of the database fields (however, you must specify date and time fields). When you specify the **-O** parameter, the ARSDOC program stores a default value in any database field that you omit. The default value for string fields is an empty string. The default value for numeric fields is `0` (zero). Numeric fields include integer and decimal (floating point) fields. When adding a document, you can specify the document data in one of three ways:

- A document file
- An SQL query that contains clauses, database field names, index values, and operators that identifies an existing document
- A public named query that identifies an existing document

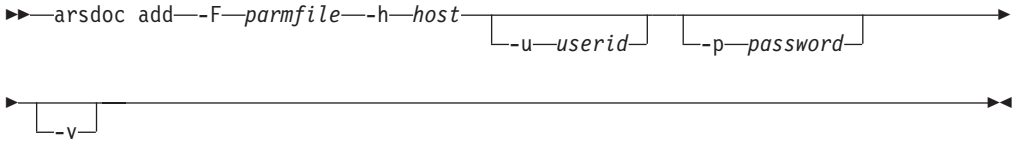
When you use the ADD function to add a row for an existing document, the row must be added to the application group that contains the document. The application group must be specified with the **-g** parameter and must be one of the application groups referenced by the folder named with the **-f** parameter.

Important: The ADD function will fail unless the Database Organization for the application group named with the **-g** parameter is Multiple Loads per Database Table and the Expiration Type is Segment or Document.

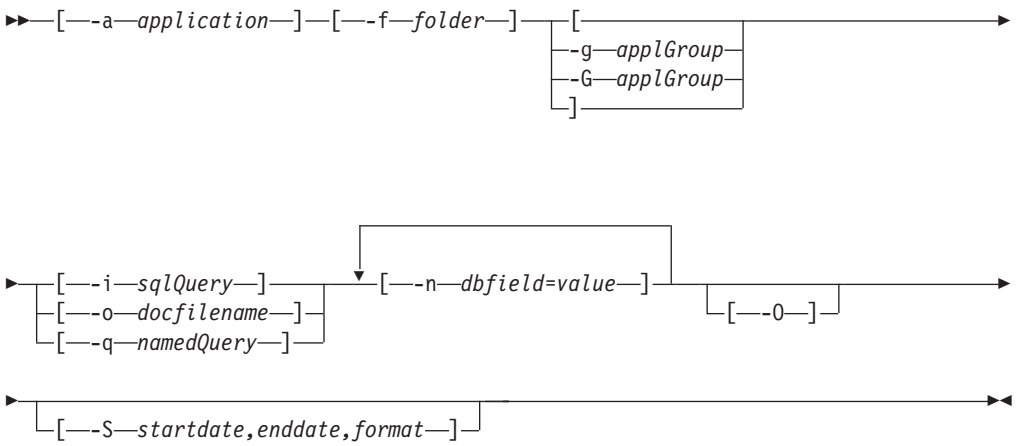




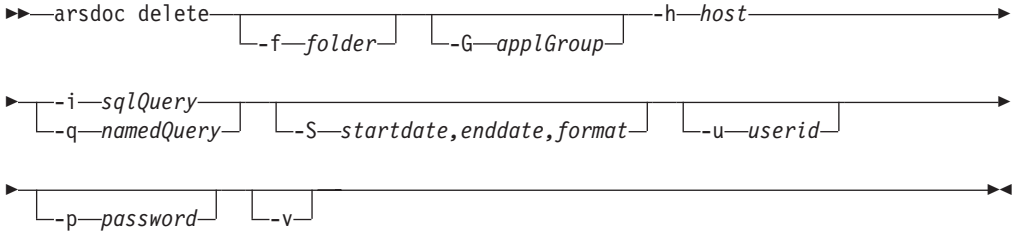
When you use the parameter file option, you must not specify the **-a**, **-f**, **-g**, **-G**, **-i**, **-n**, **-o**, **-O**, **-q**, or **-S** parameters on the command line. The following diagram shows the syntax of the ADD function when you use the parameter file option:



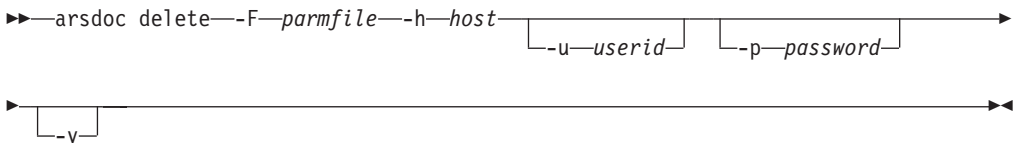
The syntax of the parameter file for the ADD function is:



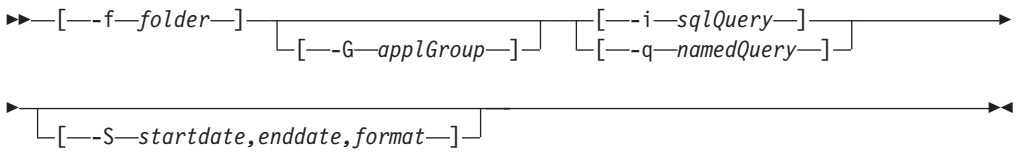
DELETE function: Use to delete documents from OnDemand. To identify the documents that you want to delete, you must enter an SQL query or specify the name of a public named query. The SQL query must contain clauses, database field names, index values, and operators. The DELETE function deletes all documents that match the query.



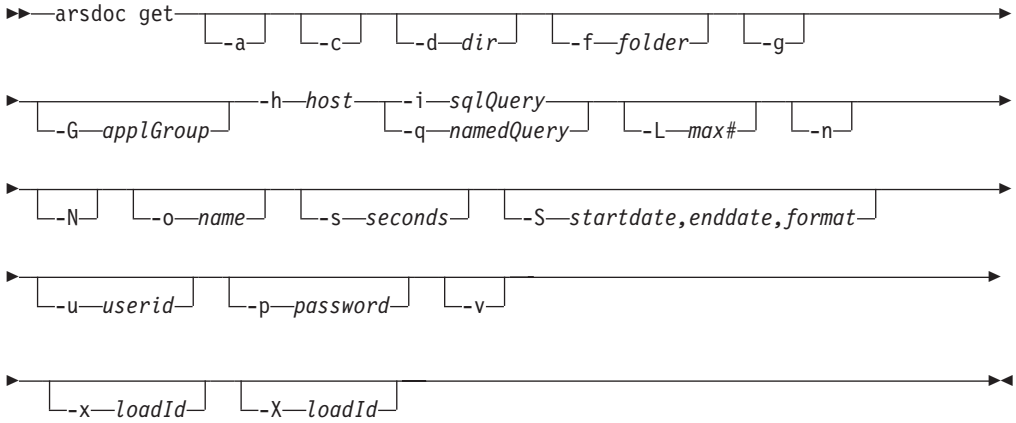
When you use the parameter file option, you must not specify the **-f**, **-G**, **-i**, **-q**, or **-S** parameters on the command line. The following diagram shows the syntax of the DELETE function when you use the parameter file option:



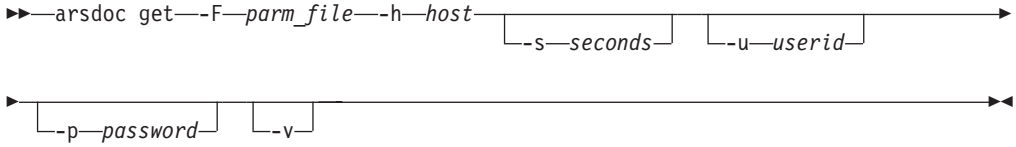
The syntax of the parameter file for the DELETE function is:



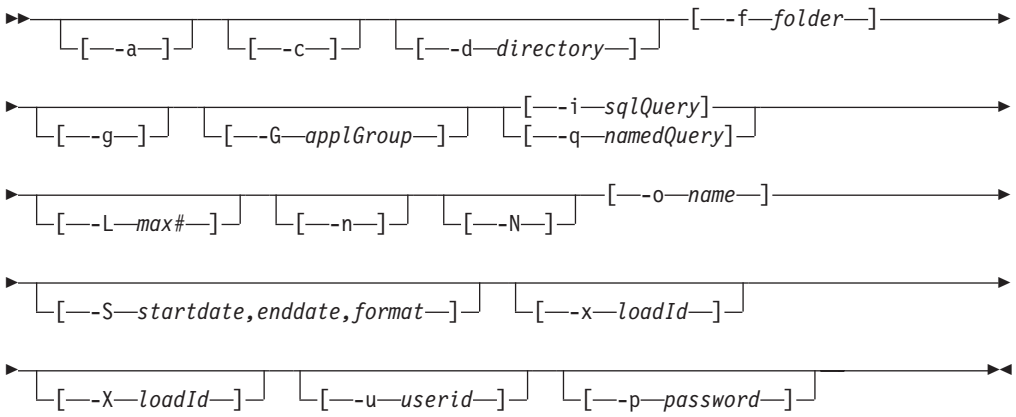
GET function: Use to retrieve documents and resources from OnDemand. The GET function can also generate and save generic index data for the documents that match the query. You must identify the name of the OnDemand server, name the folder that you want to search, and specify one or more queries to run. To query the database, you can enter an SQL query or specify the name of a public named query. The SQL query must contain clauses, database field names, index values, and operators. By default, the ARSDOC program sends a copy of the documents that match the query to the display (interactive) or the joblog (batch). You typically choose to write the output to a file.



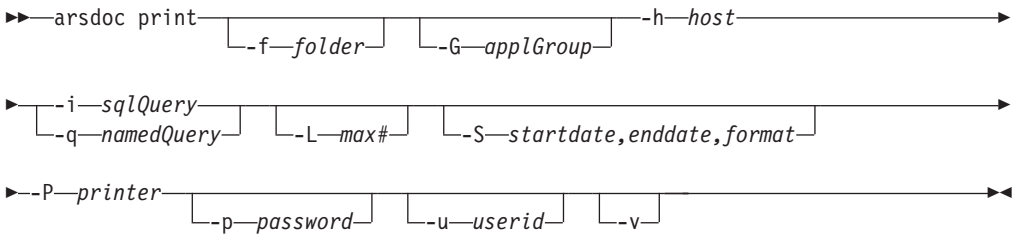
When you use the parameter file option, you must not specify the **-a**, **-c**, **-d**, **-f**, **-g**, **-G**, **-i**, **-L**, **-n**, **-N**, **-o**, **-q**, **-S**, **-x**, or **-X** parameters on the command line. The following diagram shows the syntax of the GET function when you use the parameter file option:



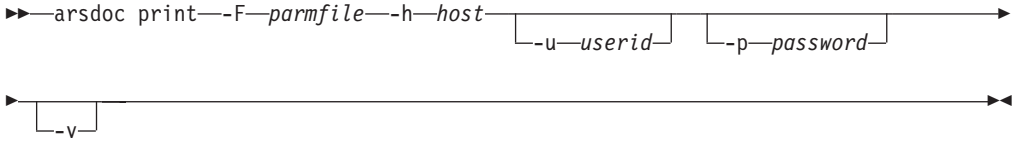
The syntax of the parameter file for the GET function is:



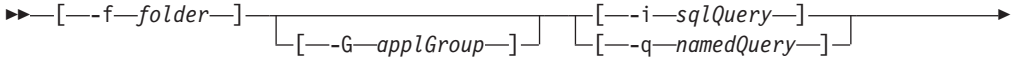
PRINT function: Use to send documents to an OnDemand server printer. You must name the OnDemand server and the folder that you want to search and specify the query to run. The items that match the query are sent to the printer named with the **-P** parameter. The printer must be defined to OnDemand using the administrative client. To query the database, you can enter an SQL query or specify the name of a public named query. The SQL query must contain clauses, database field names, index values, and operators. You can limit the number of documents sent to the printer by using the **-L** parameter and specifying the maximum number of documents that should be retrieved, regardless of the number of documents that match the query. You can limit the number of database tables searched, and possibly increase the performance of a query, by specifying the **-S** parameter and specifying a start date and an end date. The PRINT function does not currently support server fax.

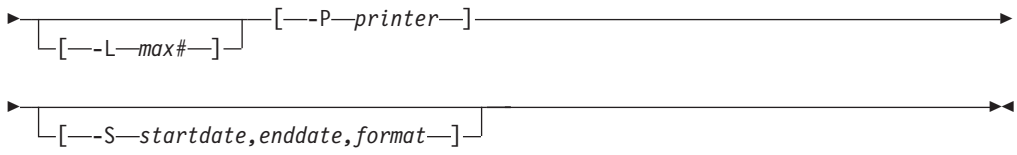


When you use the parameter file option, you must not specify the **-f**, **-G**, **-i**, **-L**, **-P**, **-q**, or **-S** parameters on the command line. The following diagram shows the syntax of the PRINT function when you use the parameter file option:

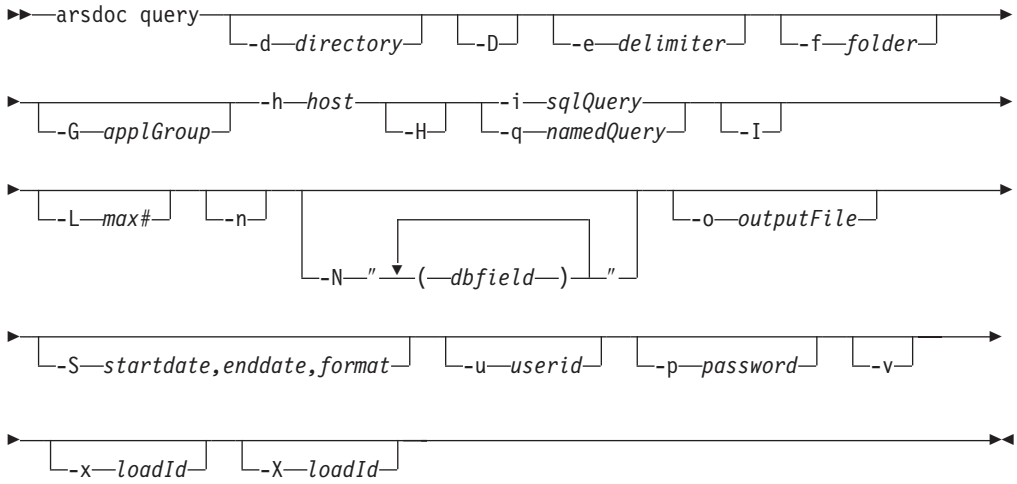


The syntax of the parameter file for the PRINT function is:

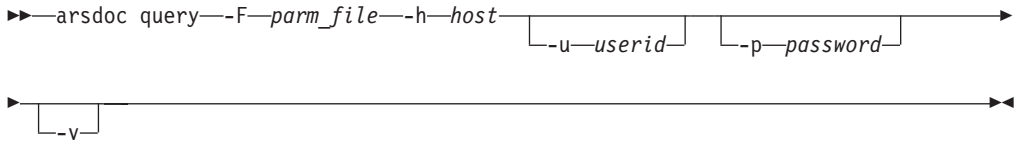




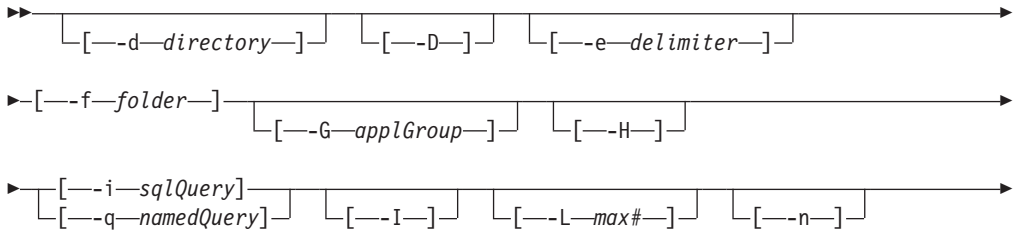
QUERY function: Use to generate a list of items that match a query. You must name the OnDemand server and the folder that you want to search and specify the query to run. To query the database, you can enter an SQL query or specify the name of a public named query. The SQL query must contain clauses, database field names, index values, and operators. By default, the ARSDOC program sends the list of items that match the SQL query to the display (interactive) or the joblog (batch). You typically choose to write the output to a file.

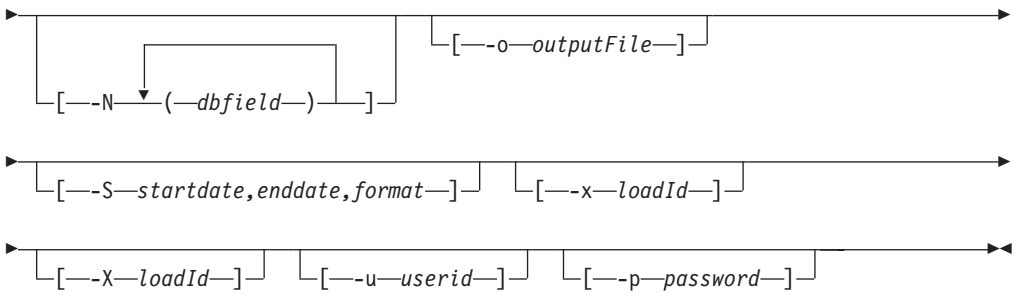


When you use the parameter file option, you must not specify the **-d**, **-D**, **-e**, **-f**, **-G**, **-H**, **-i**, **-I**, **-L**, **-n**, **-N**, **-o**, **-q**, **-S**, **-x**, or **-X** parameters on the command line. The following shows the syntax of the QUERY function when you use the parameter file option:

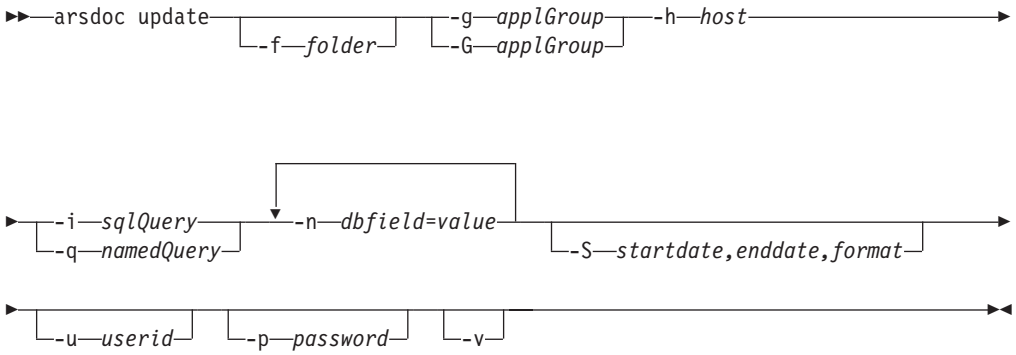


The syntax of the parameter file for the QUERY function is:

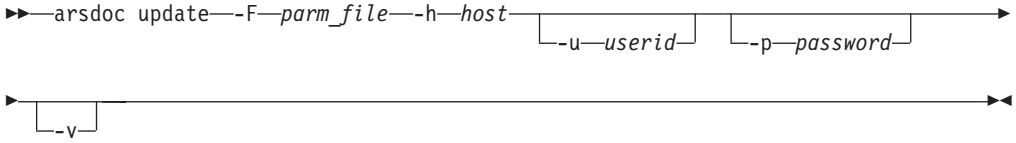




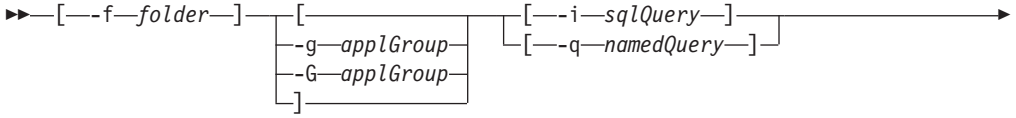
UPDATE function: Use to update documents. You must name the application group to update and specify one or more application group fields and their values. To identify the documents that you want to update, you must enter an SQL query or specify the name of a public named query. The SQL query must contain clauses, database field names, index values, and operators. The UPDATE function updates all of the rows that match the query. It does not change the actual document data.

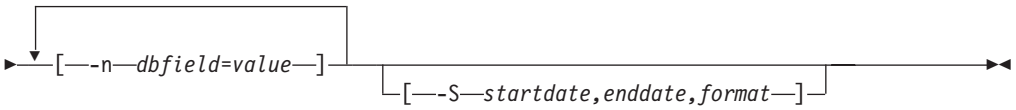


When you use the parameter file option, you must not specify the **-f**, **-g**, **-G**, **-i**, **-n**, **-q**, or **-S** parameters on the command line. The following shows the syntax of the UPDATE function when you use the parameter file option:



The syntax of the parameter file for the UPDATE function is:





Parameters

-a For the GET function, when retrieving AFP documents, specify this parameter to include resources with the documents that are retrieved. If documents from the same application have different resource groups, then the ARSDOC program creates separate output files for each resource group.

For the ADD function, you must specify the name of the OnDemand application. The application must belong to the application group named with the **-g** (or **-G**) parameter.

-c For the GET function, use to concatenate all of the AFP documents that match the query in one output file. Name the output file with the **-o** parameter. However, even if you do specify the **-c** parameter, the ARSDOC program creates separate output files when any of the following conditions occur:

- If more than one application group is referenced by the folder. The ARSDOC program creates one output file for each application group that contains items that match the query.
- If more than one application is contained in an application group. The ARSDOC program creates one output file for each application that contains items that match the query.
- If documents from the same application have different resource groups, the ARSDOC program creates separate output files for each resource group.

For example, if a folder references two application groups, then the following specification:

```
-o student -c
```

Can result in file names such as:

```
student.516
student.517
```

Where 516 and 517 are application group identifiers. One file is created for each application group. Each file contains all of the items that match the query for that particular application group.

-d directory The name of the directory where the ARSDOC program writes the output files. The directory must exist before the ARSDOC program attempts to save the output files.

-D For the QUERY function, appends the document handle information to the end of each output record. The document handle information consists of the following ten values, in the order listed:

Document name

Offset

Length

Compressed object offset

Compressed object length
Annotation type
Compression type
Resource ID
Primary node ID
Secondary node ID

The values are separated by a delimiter. The default delimiter is the comma character. You can specify a different delimiter with the **-e** parameter.

You cannot specify the **-H**, **-n**, or **-N** parameters when you specify the **-D** parameter.

-e delimiter For the QUERY function, determines the character that separates the database field values in the output. By default, OnDemand separates values in the output with a comma.

-f folder The name of the OnDemand folder. The folder name must be specified exactly as it appears in OnDemand. The case of the folder name is significant. For example, to query the OnDemand System Log folder, you must enter:

```
-f "System Log"
```

If you are using a parameter file, then you must specify the **-f** parameter in the parameter file. If you are not using a parameter file and you do not specify the **-f** parameter, then the ARSDOC program prompts you for the folder name when you run the program.

Note: A folder can be used to search one or more application groups. Because the ARSDOC program generates a single SQL query to search all of the application groups, the properties of the database fields must be the same for each application group. The properties include the field name, type, and length. For example, suppose that you define the following application groups and fields:

Application Group	Field Names
Student Bills	name, account, billDate
Student Grades	name, account, gradeDate
Student Transcripts	name, account, transcriptDate

You cannot query the application groups using the ARSDOC program because the name of the date field is not the same for each application group. However, if you were to define the application groups and fields as follows:

Application Group	Field Names
Student Bills	name, account, studentDate
Student Grades	name, account, studentDate
Student Transcripts	name, account, studentDate

Then you could query the application groups using the ARSDOC program because the names of the database fields are the same for each application group.

-F parmFile Determines the name of the file that contains the actions to run

and other document and file parameters, values, and options. You typically specify this option when you want to perform multiple actions.

Parameters and values specified in the parameter file must be enclosed in brackets.

You can use the \ (backslash) character to continue the parameters of an action to two or more lines. An action (one or more input lines) can contain a maximum of 32767 characters (bytes).

A parameter file can contain blank lines and comment lines. A comment line contains the # character in the first column.

-g For the GET function, use to generate generic index data for the items that match the query. See the *IBM Content Manager OnDemand for iSeries Common Server: Indexing Reference* for details about the generic index file.

When you specify the **-g** parameter, you must specify the **-c**, **-N**, and **-o** parameters. However, you cannot specify database field names with the **-o** parameter.

The ARSDOC program uses the following convention to name the output files that are generated with the **-g** parameter:

-o.res_id.appl_group.appl.type

Where:

-o is the value specified with the **-o** parameter

res_id is the resource group identifier. A resource group identifier always starts at 1 (one) and is incremented by one for each version of the resource group.

appl_group is the name of the application group

appl is the name of the application

type is the file type:

- **out** identifies a document file
- **ind** identifies a generic index file
- **res** identifies a resource file

In general, the number of files generated is dependent on the number of application groups in a folder, the number of applications in an application group, and the number of versions of resource groups in an application.

For the ADD and UPDATE functions, determines the name of the OnDemand application group. The application group that you specify must be able to be searched from the folder named with the **-f** parameter.

-G applGroup Use to specify the name of the application group.

For ADD: When the database query is run to retrieve the document that contains the data that is to be used in the add function, the search is limited to the specified application group, even if the folder named with the **-f** parameter can be used to search more than one application group. This ensures that only documents in the specified application group can be used for the add function. You can specify the name of the application group with the **-g** parameter or the **-G** parameter.

For DELETE: The **-G** parameter is an optional parameter. If specified, then the database query that is run to determine the document(s) to delete is limited to the specified application group. The addition of the **-G** parameter allows you to delete documents from a specific application group in folders that can search more than one application group. If you do not specify the **-G** parameter, then the query runs against all of the application groups that can be searched from the folder.

For GET: The **-G** parameter is an optional parameter. If specified, then the database query that is run to retrieve documents is limited to the specified application group. The addition of the **-G** parameter allows you to retrieve documents from a specific application group in folders that can search more than one application group. If you do not specify the **-G** parameter, then the query runs against all of the application groups that can be searched from the folder.

For PRINT: The **-G** parameter is an optional parameter. If specified, then the database query that is run to determine the document(s) to print is limited to the specified application group. The addition of the **-G** parameter allows you to print documents from a specific application group in folders that can search more than one application group. If you do not specify the **-G** parameter, then the query runs against all of the application groups that can be searched from the folder.

For QUERY: The **-G** parameter is an optional parameter. If specified, then the database query is run against the specified application group. The addition of the **-G** parameter allows you to query documents from a specific application group in folders that can search more than one application group. If you do not specify the **-G** parameter, then the query runs against all of the application groups that can be searched from the folder.

For UPDATE: When the database query is run to determine the document(s) to update, the search is limited to the specified application group, even if the folder named with the **-f** parameter can search more than one application group. This guarantees that only documents in the specified application group can be updated. You can specify the name of the application group with the **-g** parameter or the **-G** parameter.

For all functions: You can use the **-G** parameter with the **-i** parameter to query folders that can search more than one application group. For example, a folder contains three application groups; you want to query only one of the application groups. Use the **-G** parameter to specify the name of the application group that you want to query. Use the **-i** parameter to specify the application group's database field names. You can also use the **-G** and **-i** parameters when the application groups have different database field names. The following example shows how to search a folder and three application groups that have different database field names:

```
arsdoc get -f "Student Information" -G loans
    -i "WHERE number LIKE '123456' AND loanDate = 10593"
arsdoc get -f "Student Information" -G grades
```

```
    -i "WHERE number LIKE '123456' AND gradeDate = 10593"
arsdoc get -f "Student Information" -G transcripts
    -i "WHERE number LIKE '123456' AND transDate = 10593"
```

You can use the **-G** parameter with the **-q** parameter to query folders that can search more than one application group. When you specify the **-G** parameter and you specify a public named query with the **-q** parameter, the ARSDOC program queries the application group named with the **-G** parameter instead of the application group specified in the named query. (If you do not specify the **-G** parameter, then the query runs against the application group specified in the named query. If the named query does not identify an application group, then the query runs against all of the application groups that can be searched from the folder.)

-h instance The instance name. This is a required parameter.

-H For the QUERY function, specify this parameter to generate a header record in the output. The default header record contains the application group field names. By default, the field names are delimited with the comma character. You can specify a delimiter of your choice with the **-e** parameter.

You can use the **-H** parameter to generate output that contains only the application group database field names. To do so, specify the **-H** parameter without specifying the **-i** or **-q** parameters. (You also must not specify the **-L**, **-n** or **-N** parameters.) The ARSDOC program writes the database field names to the specified output file, or to the display (interactive) or the joblog (batch).

-i sqlQuery A valid SQL query, that includes the names of one or more application group database fields, index values, and operators. OnDemand does not validate the string that you enter. Please see the *DB2 Universal Database™ for iSeries: SQL Reference* for an overview of SQL concepts and details about how to construct a query.

If you specify the **-i** parameter, then you must not specify the **-q** parameter

For the DELETE or UPDATE functions, if more than one document meets the search criteria, then multiple documents will be deleted or updated. For an update, all of the documents will be updated with the same values.

To construct a query with a database field of type date, you must use the OnDemand internal format of the date. That is, the number of days since January 1, 1970. You can use the ARSDATE program to list the internal format for a given date string. For example, the following shows how to use the ARSDATE program to obtain the internal date for July 21, 1995:

```
arsdate -a 7/21/95
```

The ARSDATE program displays:

```
7/21/95 -> 9333
```

You would then enter 9333 as the index value for the date field.

-I For the QUERY function, appends the Load ID to the end of each

output record. The Load ID is separated from the database field values by a delimiter. The default delimiter is the comma character. You can specify a different delimiter with the **-e** parameter. You cannot specify the **-H**, **-n**, or **-N** parameters when you specify the **-I** parameter.

To use the **-I** parameter, the user running the query must have permission to access the System Log application group and folder.

If the Load ID is not found in the system log, then the string `Load ID could not be found` is appended to the end of the output record.

Note: The Load ID for a document is determined by searching the system log. Searching the system log can be very time consuming, depending on the number of records stored in the system log. The system log is searched for each document that matches the query.

-L max#

For the GET and PRINT functions, determines the maximum number of items retrieved from OnDemand, regardless of the number of items that match the query. OnDemand retrieves items in the order that they were loaded in the database.

For the QUERY function, determines the number of items included in the hit list, regardless of the number of items that match the query. OnDemand lists the items in the order that they were loaded into the database.

-n

For the GET function, specify to retrieve items one at a time, so that the items can be retrieved in sorted order. The items are sorted using the Sort field from the Field Information page for the folder, if specified. If the fields are not sorted (that is, `Sort=0` is specified for **all** of the fields), then the items will not be retrieved in sorted order. If you do not specify the **-n** parameter, then the GET function ignores the Sort field and you cannot predict the order in which the items will be retrieved.

For the QUERY function, use to number the items in the output file. If you specify this option, the QUERY function will sequentially number each line in the output file, beginning with `1` (one). **Note:** The QUERY function sorts the items using the Sort field from the Field Information page for the folder, if specified.

For the ADD and UPDATE functions, use to specify the application group database field names and their values using the form **-n** *dbfield*=*value*. You can specify one or more field names and their values (by specifying the **-n** parameter one time for each database field name and its value). When adding a document, you must specify all of the application group fields unless you specify the **-O** parameter. When updating a document, you can specify one or more fields and their values. For a date field, you must specify the value using the Display Format from the Field Information page under folders.

Note: Specify a blank field value by using single quotes within double quotes. For example: `-n middle="''"`. Specify a string field value that contains a blank or other special

character by enclosing the field value in single quotes within double quotes. For example: `-n name="'Sally Smith'"`.

-N For the GET function, when querying a folder that searches more than one application group or a folder that searches an application group that contains more than one application, specify this parameter to add the resource identifier, application group name, and application name to the output file name. When you specify the **-N** parameter, you must specify the **-c** parameter. If you specify the **-g** parameter to generate generic index data, you must specify the **-N** parameter.

If the folder searches more than one application group or an application group contains more than one application and you do not specify the **-N** parameter, then the ARSDOC program adds the application group or application identifier to the output file name. For example, the following specification:

```
-o student -c
```

Can result in output file names such as:

```
student.516
student.517
```

Where `516` and `517` are application group identifiers. However, when you specify the **-N** parameter, the ARSDOC program uses the resource identifier, application group name, and application name to name the output file. For example, the following specification:

```
-o student -c -N
```

Can result in output file names such as:

```
student.1.BILLS.1995
student.1.BILLS.1996
```

Where `1` is the resource identifier, `BILLS` is the application group name, and `1995` and `1996` are application names.

The number of index files created is dependent on the number of application groups in a folder, the number of applications in an application group, and the number of resource groups in an application.

For the QUERY function, determines the application group fields that the ARSDOC program writes to the output file and the field names that appear in the header record. By default, the ARSDOC program writes all fields to the output file. You can specify one or more application group field names using the form **-N(***dbfield***)**...**(***dbfield***)**. Each field name that you specify must be delimited with parenthesis. When you run a query from the command line, you must delimit the entire string in double quote characters. For example, **-N"(***dbfield1***)(***dbfield2***)"**.

-o name For the GET function, use to write documents to one or more files and identify a user-defined string used to generate unique file names. For example, the following specification:

```
-o student -c
```

Can result in the following output file name:

```
student
```

You can concatenate one or more of the database field names that you specify with the **-i** parameter to generate a unique file name. For example, the following specification:

```
-o (sdate)(student)
-i "WHERE sdate='971025' AND student='001200340056'"
```

Can result in the file name:

```
971025.001200340056
```

When you use database field names to generate a unique file name:

- OnDemand verifies that the field names that you specify are valid for the application groups that can be searched by the folder specified with the **-f** parameter.
- If the field name that you specify is a date field, the output format of the date is determined by the Format on the Load Information page under applications.
- The field names must be delimited with parenthesis.
- You can specify the fields in any order. The order that you specify determines the file name that the ARSDOC program generates.
- You cannot use a field name to represent a directory name. For example:

  ```
  -o (field_1)/(field_2)
  ```

  Is not valid.
- You cannot specify the **-c** parameter to concatenate items in one output file. Each item that matches the query is stored in a separate output file.

If more than one item matches a query and you do not generate a unique file name using database field names, concatenate items in a single file with the **-c** parameter, or specify the **-g** parameter, then the ARSDOC program generates a unique file name for each item that matches the query by adding a *.n* extension to the file name. Where *n* is the number of the item that matched the query. For example, if you specify:

```
-o statements
```

And two items match the query, the ARSDOC program creates the following files:

```
statements.1
statements.2
```

You must specify the **-o** parameter when you specify the **-c** parameter.

For the QUERY function, determines the file name of the output file in which the ARSDOC program writes the list of items that match the query.

For the ADD function, determines the name of the input file that contains the document to be added. The value that you specify is not checked for valid characters. You can specify a full path name, including the forward slash characters as part of a directory path. When adding a document, you can provide the input data by specifying the name of the input file that contains the data with the **-o** parameter, an SQL query with the **-i** parameter, or a public named query with the **-q** parameter. Only one document can be added at a time.

-O
For the ADD function, you must specify this parameter if you intend to omit one or more database fields. However, you can never omit date or time fields. When you specify the **-O** parameter, the ARSDOC program stores a default value in any other database field that you omit. The default value for string fields is an empty string. The default value for numeric fields is `0` (zero). Numeric fields include integer and decimal fields.

-p password
The password of the OnDemand user that you named with the **-u** parameter. If there is no password assigned to the user that you specify, use double quotes to show a null password. That is, specify **-p** `""`. If you omit the **-p** parameter and do specify a userid with the **-u** parameter, then the ARSDOC program prompts you to enter the password when you run the program. If there is no password assigned to the user that you specify, press the Enter key when prompted. If you omit **-u** and **-p** parameters, OnDemand will use the OS/400 user profile of the user issuing ARSDOC as the OnDemand userid.

-P printer
For the PRINT function, identifies the OnDemand server printer to which you want to send the documents that match the query.

-q namedQuery
The name of a public named query for the folder named with the **-f** parameter. A named query is a set of search criteria previously saved on the server that can be recalled by name to search a folder. A named query is typically defined to search a folder for a specific document or set of documents.

If you specify the **-q** parameter, you cannot specify the **-i** parameter or the **-S** parameter.

-s seconds
For the GET function, determines the number of seconds that the ARSDOC program waits between query requests when you specify more than one query with the **-F** parameter. If you do not specify this option, then the ARSDOC program does not wait between query requests. That is, the default is `0` (zero) seconds.

-S startdate,enddate[,format]
Provides a date range that the ARSDOC program uses to limit a search to specific tables. When you specify this parameter, the ARSDOC program searches only tables that contain a segment within the specified date range.

An optional date format can also be specified with the **-S** parameter. For example: `-S "01011990,12311990,%m%d%Y"`, where `%m%d%Y` is the date format. See "ARSDATE" on page 149 for a list of the standard date formats. If you do not specify a date format,

then the date values that you specify must be in the date format that was selected for the Display Format field on the Field Information page for the folder.

Note: For most queries, you should always specify this parameter and a date range. Doing so limits the range of a query and can significantly improve the performance of a query. For the ADD function, if you specify the **-o** parameter, you cannot specify the **-S** parameter. For all functions, if you specify the **-q** parameter, you cannot specify the **-S** parameter.

-u userid The OnDemand user that is permitted to perform the specified function. The ARSDOC program verifies that the userid that you specify is a valid OnDemand user for the server that you name with the **-h** parameter, that the userid is permitted to open the folder that you name with the **-f** parameter, and that the userid has permission in application groups to perform the specified function. If you omit the **-u** parameter, the current user profile of the job running ARSDOC is used as the OnDemand userid. If that user profile does not match an OnDemand userid, then ARSDOC prompts you to enter the userid.

-v Enables verbose mode, which displays all messages (informational and error). By default, the ARSDOC program displays error messages.

-x loadId For the GET function, specify to limit the documents that can be retrieved to the set of documents that were loaded into the system under the specified *loadId*.

For the QUERY function, specify to limit the query to the set of documents that were loaded into the system under the specified *loadId*.

When you specify the **-x** parameter, use the **-f** and **-G** parameters as follows:

- Specify the **-f** parameter to search all application groups. You can specify the search using the **-i** parameter or the **-q** parameter.
- Specify the **-G** parameter to search a specific application group. You must specify the search using the **-i** parameter.
- Specify both the **-f** parameter and the **-G** parameter. The ARSDOC program will verify that the application group can be searched from the folder.

When you specify the **-x** parameter, you cannot specify the **-X** parameter.

The following example shows how to limit a query to the documents that were loaded into the system under a specific load ID. For example, assume that the specified folder could be used to search several application groups; each application group contains more than one application; there are 500,000 documents in the application groups. By specifying the **-x** parameter, the query will be limited to the set of documents that was loaded into the system under the specified load ID; without the **-x** parameter, the query is run against all 500,000 documents.

```
arsdoc query -h rockies -x 19867-025-0-3FAA-10136-10136
  -f load-scanned-images -q query -o output -v
```

-X loadId For the GET function, specify the **-X** parameter to retrieve documents by using the index file that was generated for the specified *loadId*.

For the QUERY function, specify the **-X** parameter to build a hit list from the index file that was generated for the specified *loadId*.

When you specify the **-X** parameter, you must specify the **-G** parameter and name the application group. When you specify the **-X** parameter, you cannot specify the **-x** parameter, or the **-i**, **-q**, **-S**, and **-f** parameters.

The following example shows how to use an index file to retrieve documents from the system. The index file was generated for the set of documents that was loaded into the system under the specified load ID. When using the **-X** parameter, the database is not queried; rather, documents are retrieved based on the information in the index file.

```
arsdoc get -h rockies -X 19867-025-0-3FAA-10136-10136
-G load-scanned-images-jpeg -o jpeg -v
```

Examples

1. The following shows how to run the ARSDOC program using the command line parameters to retrieve an item from OnDemand and save a copy of the item in a file in the current directory.

   ```
   arsdoc get -h rockies -f "Student Information" -o student
    -c -S 1/1/97,12/31/97 -i "WHERE student='001200340056'" -v
   ```

   The ARSDOC program concatenates all items that match the query in the following output file:

   ```
   student
   ```

2. The following shows how to run the ARSDOC program and specify the name of a parameter file in the current directory that contains a set of queries. The results of each query are saved in separate files in the current directory.

   ```
   arsdoc get -h rockies -F parmfile -v
   ```

   The parameter file contains two queries:

   ```
   [-f "Student Information"] [-S 1/1/97,12/31/97] [-o (student)(type)] \
   [-i "WHERE student='123420010056' AND type='B' OR type='G' OR type='T'"]

   [-f "Student Information"] [-S 1/1/97,12/31/97] [-o (student)(type)] \
   [-i "WHERE student='123450011917' AND type='B' OR type='G' OR type='T'"]
   ```

   Assuming that items exist in the database for all of the specified types, the ARSDOC program creates the following output files:

   ```
   123420010056.Bills
   123420010056.Grades
   123420010056.Transcripts
   123450011917.Bills
   123450011917.Grades
   123450011917.Transcripts
   ```

3. The following shows how to run the ARSDOC program using the command line parameters to retrieve items from OnDemand, write the documents to a file, and generate and save generic index data for the items that match the query. The name of a public named query is identified.

```
arsdoc get -h rockies -f "Student Information"
 -a -c -g -o student -q "3rd yr students GPA>3.5" -N -v
```

The number of output files generated is a factor of the number of application groups queried, the applications contained in the application groups, whether the data is AFP, and if so, the versions of resource groups in each application. At a minimum, for AFP data with one version of the resource group and one application group, the ARSDOC program generates three output files. For example:

```
student.1.Student Information.TRANSCRIPTS.ind
student.1.Student Information.TRANSCRIPTS.out
student.1.Student Information.TRANSCRIPTS.res
```

For AFP data, if there is more than one version of the resource group per application, then the ARSDOC program can generate additional output files. For example:

```
student.1.Student Information.TRANSCRIPTS.ind
student.1.Student Information.TRANSCRIPTS.out
student.1.Student Information.TRANSCRIPTS.res
student.2.Student Information.TRANSCRIPTS.out
student.2.Student Information.TRANSCRIPTS.ind
student.2.Student Information.TRANSCRIPTS.res
```

If the application group contains more than one application, then the ARSDOC program can generate additional output files. For example:

```
student.1.Student Information.BILLS.ind
student.1.Student Information.BILLS.out
student.1.Student Information.BILLS.res
student.1.Student Information.GRADES.ind
student.1.Student Information.GRADES.out
student.1.Student Information.GRADES.res
```

4. The following shows how to run the ARSDOC program using the command line parameters to generate a list of items and save the list in a file in the current directory.

```
arsdoc query -h rockies -f "Student Information"
 -o query1.out -H -S 1/1/97,12/31/97
 -i "WHERE student='0012-0034-0056' AND
     type='B' OR type='G' OR type='T'" -v
```

5. The following shows how to run the ARSDOC program and specify the name of a parameter file in the current directory that contains a set of queries. The result of each query is saved in a separate file in the current directory.

```
arsdoc query -h rockies -F parmfile -v
```

The parameter file contains three queries:

```
[-f "Student Information"] [-i "WHERE type='B'"] [-o "query2.out"] \
[-S 1/1/97,12/31/97] [-H] [-N (student)(id)(p_date)]

[-f "Student Information"] [-i "WHERE type='G'"] [-o "query3.out"] \
[-S 1/1/97,12/31/97] [-H] [-N (student)(id)(p_date)]

[-f "Student Information"] [-i "WHERE type='T'"] [-o "query4.out"] \
[-S 1/1/97,12/31/97] [-H] [-N (student)(id)(p_date)]
```

6. The following example shows how to use the ARSDOC program to add a document to the Credit Card Statements folder:

```
arsdoc add -h rockies -o/newdata/crd.dat
  -n "crd_date=01/21/98" -n "account='000-000-000'"
  -n balance=123.45 -n "name='John Watpole'"
  -f "Credit Card Statements" -g CRD -a CRD -v
```

7. The following example shows how to use the ARSDOC program to update a document in the Credit Card Statements folder, changing the balance from 123.45 to 0.00:

```
arsdoc update -h rockies
  -i "where account='000-000-000' and name='John Watpole'"
  -n "balance=0.00" -S 1/1/97,12/31/97
  -f "Credit Card Statements" -g CRD -v
```

8. The following example shows how to use the ARSDOC program to delete a document from the Credit Card Statements folder:

```
arsdoc delete -h rockies
  -i "where account='000-000-000' and name='John Watpole'"
  -f "Credit Card Statements" -S 1/1/97,12/31/97 -v
```

9. The following example shows how to use the ARSDOC program to send items that match a query to the OnDemand server print function:

```
arsdoc print -h rockies -P svrprt1
  -i "where account='000-000-000' and name='John Watpole'"
  -f "Credit Card Statements" -S 1/1/97,12/31/97 -v
```

10. You can use the ARSDOC program to add an index that points to an existing document. For example, assume that you loaded a bank statement for account number 000-000-000, date 5/23/97, and account name Joe Smith. You want to add a new index, but point to the existing statement. The new index uses the same account number and date, but contains a different account name (for example, Sally Smith). After adding the index, if a query is run with account name Joe Smith or Sally Smith, the same bank statement will be retrieved. To add an index for an existing document:

```
arsdoc add -h rockies
  -i "where sdate=10005 and account='000-000-000' and name='Joe Smith'"
  -n "sdate=5/23/97" -n "account=000-000-000" -n "name='Sally Smith'"
  -f "Credit Card Statements" -S 1/1/97,6/31/97 -v
```

Notes

The ADD function will fail unless the Database Organization for the application group named with the **-g** parameter is Multiple Loads per Database Table and the Expiration Type is Segment or Document.

The ADD function can be run without providing an input document from a file or by retrieving an existing document from the system. This means that you can add database field values without adding a document. To add database field values without adding a document, do not specify the **-o**, **-i**, or **-q** parameters; specify the database field names and their values using one or more **-n** parameters.

IFS location

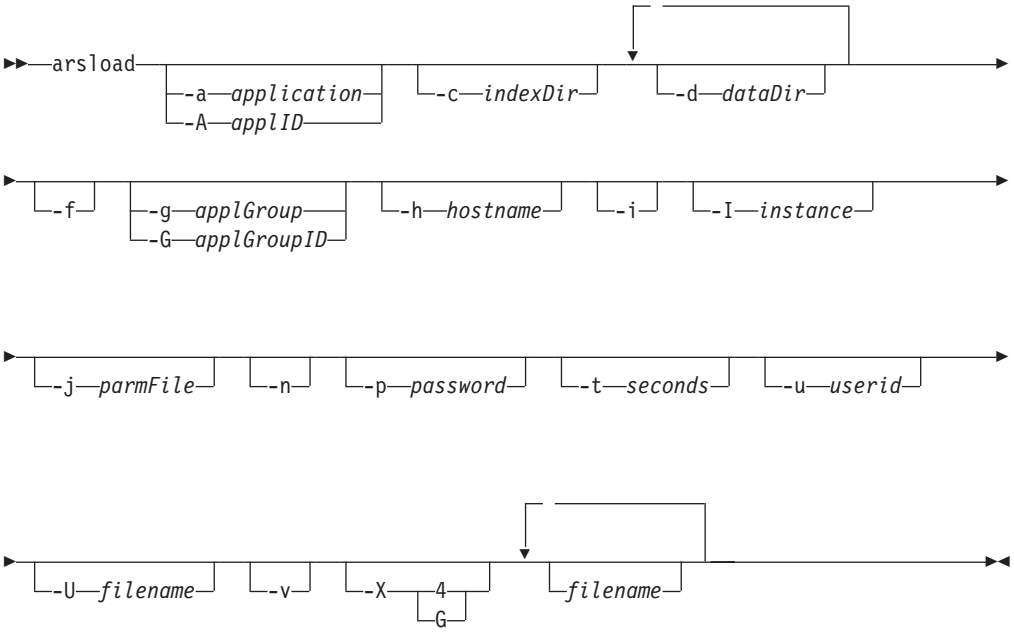
/usr/bin/arsdoc The executable program.

ARSLOAD

Purpose

The ARSLOAD program can be used to process the input files (not spooled files) that you want to load into the system. The ARSLOAD program determines if the input data needs to be indexed, and if so, calls the appropriate indexing program. The ARSLOAD program calls the storage manager programs to load report data on storage volumes and the database manager to update the OnDemand database with the index information that was extracted from or specified for the input file.

Syntax



Description

The ARSLOAD program is the main OnDemand data indexing and loading program. The ARSLOAD program calls the indexing program if the input data needs to be indexed, creates input files for the storage manager(s), and updates the OnDemand database. The ARSLOAD program saves processing messages in the system log. You can open the System Log folder and list the messages that were generated when an input file was processed.

You typically configure the ARSLOAD program to run as a monitor to periodically check specified IFS directories for input files to process.

When you run the ARSLOAD program, you must provide the userid and password of an OnDemand user with administrator authority for the application group into which the input data will be loaded. There are several ways that you can provide the userid and password:

- Do not specify the **-u** and **-p** parameters. In this case, the current user profile of the job that is running ARSLOAD is used as the OnDemand userid. This is the recommended method.
- Specify the **-u** and **-p** parameters each time that you run the ARSLOAD program.
- Specify the **-U** parameter each time that you run the ARSLOAD program to name a file that contains the userid and password of a user with administrator authority for the application group.

Parameters

-a application The name of the application to load. If the application group contains more than one application, then you must identify the application to load. Otherwise, the load will fail. If you plan to automate the loading of files into different application groups and applications, then use the **-A** parameter to specify the part of the file name that identifies the name of the application.

-A applID Determines the part of the file name that identifies the application to load. If the application group contains more than one application, then you must identify the application to load. Otherwise, the load will fail. You typically use this parameter when you run the ARSLOAD program as a monitor to automate the loading of files into different application groups and applications. For example, you can transmit files from OS/390® systems to the server. A transmitted file must use the following file naming convention:

```
MVS.JOBNAME.DATASET.FORMS.YYYYDDD.HHMMSST.ARD
```

- By default, the ARSLOAD program uses the `FORMS` part of the file name to identify the application group to load. You can use the **-G** parameter to specify a different part of the file name that identifies the application group. For example, `arsload -G JOBNAME`.
- If the application group to load contains more than one application (source of data), then you must identify the application to load. Otherwise, the load will fail. For example, to use the `DATASET` part of the file name to identify the application, run the ARSLOAD program with the `-A DATASET` parameter.

Choose one of the `MVS`, `JOBNAME`, `DATASET`, and `FORMS` parts of the file name to identify the application.

-c indexDir The directory in which OnDemand temporarily stores data created by the indexing program. The default location is the directory from which the ARSLOAD program was invoked.

-d dataDir The directory that contains input files to process. The ARSLOAD program will search for input files to load in the directory that you specify. **Note:** To specify more than one directory, specify the **-d** parameter multiple times. In the following example:

```
arsload ... -d dir1 -d dir2 -d dir3 ...
```

The ARSLOAD program will search for input files in the `dir1`, `dir2`, and `dir3` directories.

-f Use to unload the data if the load process fails. If the database manager step fails, then OnDemand should remove any index data that was added to the database. If the storage manager step fails, then OnDemand should remove any storage objects that were copied to storage volumes.

Note: If an input file fails to load, then you should review the message log that was created during the load process. You can retrieve a message log from the system log. If the message log contains a Load ID, then it means that for some reason, OnDemand stored at least some of the input data in the application group. Before you attempt to reload the input data, you must remove the data that was created during the failed load process by using the RMVRPTOND command. See Chapter 14, "Deleting a report" on page 123 for help with removing the data that is created when a file is loaded into the system.

-g applGroup The name of the application group to load. This parameter is required if you specify the **loadFilename** parameter. This parameter is optional if you specify the **-d** parameter. If you

specify the **-d** parameter, then by default, the ARSLOAD program uses the FORMS part of the file name to determine the name of the application group to load. If you plan to automate the loading of files into different application groups and applications, then you should use the **-G** parameter to specify the part of the file name that identifies the application group to load.

-G applGroupID

Determines the part of the file name used to identify the name of the application group to load. You typically use this parameter when you run the ARSLOAD program as a monitor to automate the loading of input files into different application groups and applications. You can transmit files from OS/390 systems to the server. A transmitted file must use the following file naming convention:

```
HLQ.JOBNAME.DATASET.FORMS.YYYYDDD.HHMMSST.ARD
```

- By default, the ARSLOAD program uses the FORMS part of the file name to identify the application group to load. You can use the **-G** parameter to specify a different part of the file name that identifies the application group. For example, `arsload -G JOBNAME`.
- If the application group to load contains more than one application (source of data), then you must identify the application to load. Otherwise, the load will fail. For example, to use the DATASET part of the file name to identify the application, run the ARSLOAD program with the `-A DATASET` parameter.

Choose one of the MVS, JOBNAME, DATASET, and FORMS parts of the filename to identify the application group.

-h hostname The TCP/IP address, fully-qualified host name, or host name alias of the iSeries system on which the OnDemand server is running.

-i Use to run the data indexing program only; do not copy report data to storage volumes or add the index data to the database.

-I instance The name of the OnDemand instance that you want to process. The name of the default instance on the iSeries is QUSROND. You must specify the **-I** parameter and name the instance, if you are running more than one instance on the same iSeries system and you want to process an instance other than the default instance (QUSROND).

-j parmFile Use to specify the name of a file that contains additional indexing parameters.

When you specify the **-j** parameter and name a file, the ARSLOAD program adds the indexing parameters from the specified file to the indexing parameters that it extracts from the application. (Indexing parameters are typically specified on the Indexing Information page in applications.) If an indexing parameter appears in both the application and the file that you specify, then the ARSLOAD program uses the value from the file.

-n Determines whether OnDemand deletes the input files when the ARSLOAD program ends.

- If you specify the **-n** parameter, then OnDemand does not delete the input files when the ARSLOAD program ends.

- If you do not specify the **-n** parameter, then OnDemand deletes the input files when the ARSLOAD program ends.

In either case, if the ARSLOAD program fails in the load step because of a device or system problem, then you can restart the load step after you correct the problem by using the intermediate files that were created by the ARSLOAD program. The ARSLOAD program stores the intermediate files in the directory named with the **-c** parameter (or the current directory, if you did not specify the **-c** parameter). The intermediate files have the same file name as the original input file.

-p password The password for the user specified with the **-u** parameter. If the user is not assigned a password, enter a null password (that is, specify `-p ""`).

-t seconds Determines the polling time in seconds. This is the interval of time in which the ARSLOAD program checks the input directories for input files to process. The default value is 600 seconds, which means that the ARSLOAD program checks the input directories every ten minutes.

-u userid The userid of an OnDemand user with administrator authority for the application group. The user must have permission to add documents to the application group.

If you omit the **-u** and **-p** parameters, then the current user profile for the job is used for the OnDemand userid. If that user profile does not match an OnDemand userid, then OnDemand attempts to obtain the userid and password from the value of the **-U** parameter. If the ARSLOAD program cannot locate or use the file named with the **-U** parameter, then it attempts to log on to OnDemand and access the application group with a userid of `QONDADM` and no password.

-U filename See **-u** for an explanation.

Important: After you configure the file named with the **-U** parameter, remember to change the password any time that you change the user's password in OnDemand. Otherwise the load will fail. The ARSLOAD program can accept an expired password. However, the ARSLOAD program will fail if you specify an incorrect password.

-v Enables verbose mode, which displays all messages (informational and error). By default, the ARSLOAD program displays error messages.

-X indexer Allows you to override the indexing program that was specified on the Indexer Information page for the application. The possible values are `4` and `G`. Specify `-X 4` to use the OS/400 indexer. Specify `-X G` to use the OnDemand Generic Indexer.

You must specify an indexing program if the input contains index data in a format other than the one supported by the indexing program that was specified for the application. For example, assume that you defined an application and specified the OS/400 Indexer as the indexing program. Later, you used the ARSDOC program to extract documents from an application group. The

ARSDOC program generates index data that is in the Generic Indexer format. To load the index data into the application, you must specify:

```
arsload -X G . . .
```

filename The names of the input file or files to process. (If you specify more than one input file, then separate the file names with a blank.) The ARSLOAD program concatenates the following file type extensions to the file name that you specify: IND, OUT, and RES

- The IND file contains the index data
- The OUT file contains the report data
- The RES file contains the resource data

For example, if you specify `arsload -g BILLS po3510`, then the ARSLOAD program looks for the following files:

- po3510.ind
- po3510.out
- po3510.res

If the ARSLOAD program does not find a file with the IND file type extension, it automatically calls the indexer specified in the application to index the input file.

Note: If you do not specify a file name, then the ARSLOAD program checks the directory specified with the **-d** parameter for input files to process. Any files with a file type extension of ARD or PDF will be processed.

Examples

1. The following shows how to run the ARSLOAD program to check the specified directory for input files to process. The input files must have a file type of `ARD` or `PDF`. The ARSLOAD program stores temporary work files in the specified directory. The ARSLOAD program uses the `FORMS` part of the file name to determine the application group to load, and the application group contains only one application. The ARSLOAD program logs on to OnDemand and accesses the application group with the current user profile.

   ```
   arsload -c /arsacif/acif1 -d /arsacif/acif2
   ```

2. The following shows how to run the ARSLOAD program to check the specified directory for input files to process. The input files must have a file type of `ARD` or `PDF`. The ARSLOAD program stores temporary work files in the specified directory. The ARSLOAD program uses the `JOBNAME` part of the file name to determine the application group to load and the `DATASET` part of the file name to determine the application to load. The ARSLOAD program logs on to OnDemand and accesses the application group with the current user profile.

   ```
   arsload -c /arsacif/acif1 -d /arsacif/acif2 -A DATASET -G JOBNAME
   ```

3. The following shows how to run the ARSLOAD program to load the specified file into the specified application group. The ARSLOAD program logs on to OnDemand and accesses the application group with the userid and password provided with the **-u** and **-p** parameters.

   ```
   arsload -g BILLS -u bob -p secret PO3510
   ```

4. The following shows how to run the ARSLOAD program to load several input files into the specified application group. The application group name contains an embedded blank character, and must be quoted. The ARSLOAD program logs on to OnDemand and accesses the application group with the userid and password provided with the **-u** and **-p** parameters.

```
arsload -g "ABC Credit" -u bob -p secret RW7505 RW8505
```

5. The following shows how to run the ARSLOAD program to load the specified file into the specified application group. The TCP/IP address of the OnDemand library server is specified. This format of the command can be used to load input files from an object server and store the index data on the library server. The ARSLOAD program logs on to OnDemand and accesses the application group with the userid and password from the file named with the **-U** parameter.

```
arsload -g "BILLS" -h "9.99.111.222" -U userpw.cfg ARSP0MST
```

Notes

An input file cannot exceed 4 GB in size.

The OnDemand server program must be running, otherwise the ARSLOAD program will fail.

IFS location

/usr/bin/arsload The executable program.

User exit reference

Output queue monitor user exit program

The output queue monitor function (started by the STRMONOND command or by using a monitor definition in iSeries Navigator) automatically processes spooled files from selected output queues. When a spooled file is selected from an output queue to be processed, OnDemand first checks for a user exit program in the monitor job's library list with a name that matches the spooled file attribute that the monitor chose for the application name (***SPLFNAME**, ***FORMTYPE**, ***USERDATA**, ***JOBNAME**, ***USRDFNOPT1**, ***USRDFNOPT2**, ***USRDFNOPT3**, ***USRDFNOPT4**, or ***USRDFNDTA**). If a program is found, it is called, allowing the application name or application group name to be changed as needed. Then OnDemand uses the new name as it looks for an application definition to use. If a user exit program is not found, OnDemand looks for an application name that matches the original spooled file attribute.

When OnDemand calls the user exit, it passes several parameters, only two of which should be changed. These two changeable fields are:

- **Application group** (60 bytes)
- **Application** (60 bytes)

You can design an output queue monitor user exit program to alter the application group name or application name that the monitor finds as it processes the particular spooled files in the output queue. For example, if the monitor finds PGM123 in the User Data spooled file attribute of the spooled file, but you want OnDemand to use the application name of INVOICES, you can write your user exit program as the example shows:

```
CHANGE APPNAME TO 'INVOICES'
```

There is a sample monitor output queue user exit program in source file QSAMPLES2 in libraries QRDARS and QUSRRDARS. (Any program source code that you modify should not be placed in QRDARS because that library is replaced during software upgrades. However, QUSRRDARS library is not replaced and can

be used for your modified source.) IBM recommends that you copy this source code into a backup member in case you need to go back to the original sample as shipped from IBM.

Facsimile user exit program

The facsimile user exit program is designed to enable OnDemand server fax functions to work with iSeries facsimile (fax) software other than Facsimile Support from IBM. To accomplish this, modify and recompile the QRLMSFAX program source code to change the command that is issued when an end user requests a server fax during report retrieval. Simply change the line in the program that issues the SNDFAX command to use the command for the fax software you have installed on your system. (You could also call a program if no command is available.)

An end-user requesting a server fax causes the QRLMSFAX program to be called. If you do not change and recompile it, the standard program will run (which issues the Facsimile Support command). If you do change and recompile the QRLMSFAX program, then the changed program will run (which will issue **your** facsimile command).

The sample CL source code for this program (member name QRLMSFAX) can be found in source file QSAMPLES2 in libraries QRDARS and QUSRRDARS. (Any program source code that you modify should not be placed in QRDARS because that library is replaced during software upgrades. However, QUSRRDARS library is not replaced and can be used for your modified source.) IBM recommends that you copy this source code into a backup member in case you need to go back to the original function as shipped from IBM. When you recompile your program, be sure to preserve the QRLMSFAX program name and place the compiled program back in QRDARS library to replace the program that is shipped by IBM.

Appendix E. Automating ARSLOAD data loading

This section provides information to help you configure the ARSLOAD data loading API to automatically start when ever the server is initialized.

ARSLOAD

The ARSLOAD program is the main OnDemand data loading and indexing API. You can configure the ARSLOAD API to monitor specific IFS directories for report data. If the data needs to be indexed, then the ARSLOAD program calls the indexing program that is specified in the OnDemand application. The ARSLOAD program then works with the database manager to load the index data into the database and works with the storage manager to load the report data and resources on to storage volumes.

Note: The OnDemand server program must be running, otherwise the ARSLOAD program will fail.

Automating ARSLOAD

The following shows an example of the ARSLOAD API.

```
arsload -v -c /arsacif/acif4 -d /arsacif/acif1 -d /arsacif/acif2 -d
/arsacif/acif3
```

In the example, the ARSLOAD program checks for data in the specified directories every ten minutes (the default polling time). If data needs to be indexed, then the ARSLOAD program stores the index data in the specified index directory.

You must verify the names of the directories. Replace the strings `/arsacif/acif1`, `/arsacif/acif2`, `/arsacif/acif3`, and `/arsacif/acif4` with the names of directories that are valid on the server that you are configuring.

Important: The ARSLOAD program uses a particular part of the input file name to determine the application group to load. You can use the `-G` parameter to specify a different part of the file name to identify the application group to load. If the application group contains more than one application, then you must identify the application to load. Otherwise, the load will fail. You can use the `-A` parameter to specify the part of the file name that identifies the application.

If a userid and password are not specified on the ARSLOAD program, the current OS/400 userid is used See Appendix D, “API and user exit reference” on page 145 for more information about the ARSLOAD program and its userid and password parameters.

After indexing the data, the ARSLOAD program deletes the input files, unless you specify otherwise. Any output or error messages that are generated by the ARSLOAD program are written to the System Log. You can open the System Log folder and retrieve any messages that were generated by the ARSLOAD program.

© Copyright IBM Corp. 2001, 2002

Appendix F. Alternative to starting the administrative client

Starting the Administrator

If you do not need to work with the OnDemand administrative functions that are supported directly thru the OnDemand iSeries Navigator component (such as tape volumes, tape devices, optical volumes, optical storage groups, disk pool storage groups, monitor definitions, or migration policies), you can go directly to the OnDemand Administrator by following the instructions below:

1. Click Start.
2. Select Programs, then choose IBM OnDemand32.
3. Click OnDemand32 Administrator.
4. When you start the Administrator, OnDemand opens the administrator window that contains a menu bar, toolbar, navigator pane, list pane, and status bar. Figure 44 shows the administrator window.

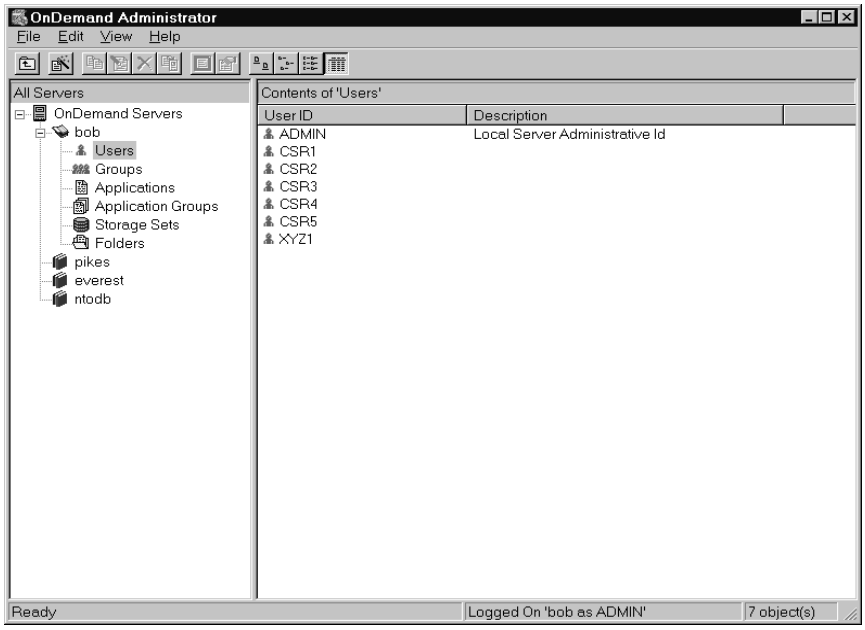


Figure 44. Administrator main window

Administrator start up parameters

OnDemand provides parameters that you can specify as properties that the operating system uses when you start the Administrator program. The parameters can be used, for example, to automate the logon process and to select the areas that appear in the navigator pane. Table 12 lists the start up parameters for the Administrator.

Table 12. Administrator start up parameters

Parameter	Purpose	Example
/1 *location*	Identifies the drive and full path name of the national language environment program file directory.	/1 D:\Program Files\IBM\OnDemand32\Locale\Enu

© Copyright IBM Corp. 2001, 2002

Table 12. Administrator start up parameters (continued)

Parameter	Purpose	Example
/B	Include Applications in the navigator pane.	/B
/D	Include Folders in the navigator pane.	/D
/D 2	Include Folders in the navigator pane, but display only the Permissions and Field Information pages.	/D 2
/E	Include Storage Sets in the navigator pane.	/E
/I	Include Users in the navigator pane.	/I
/O	Include Groups in the navigator pane.	/O
/P *password*	The password for the OnDemand user identified with the **/U** parameter. Use with the **/S** and **/U** parameters to log on to a specific server, without displaying the Logon dialog box. If the logon is not successful, then OnDemand opens the Logon dialog box to allow the user to log on to the system.	/P password
/Q	Include Application Groups in the navigator pane.	/Q
/R	Include Printers in the navigator pane.	/R
/S *server*	Identifies the logon server. Use with the **/U** and **/P** parameters to log on to a specific server, without displaying the Logon dialog box. If the logon is not successful, the OnDemand opens the Logon dialog box to allow the user to log on to the system.	/S broncos
/T	Determines the name that appears on the title bar of the administrator window. The default name is OnDemand Administrator.	/T Customer Service
/U *userid*	Identifies the OnDemand user. Use with the **/S** and **/P** parameters to log on to a specific server, without displaying the Logon dialog box. If the logon is not successful, OnDemand opens the Logon dialog box to allow the user to log on to the system.	/U admin

Changing start up parameters

To change start up parameters:

1. Click Start.
2. From the pop-up menu, select Explore.
3. Under Start Menu, expand Programs.
4. Select IBM OnDemand32.
5. In the list pane, point to OnDemand32 Administrator and click the right mouse button.
6. From the pop-up menu, select Properties.
7. Click the Shortcut tab.
8. In the Target field, enter the startup parameters. Make sure that you enter any startup parameters after the string:

```
"D:\Program Files\IBM\OnDemand32\ARSADM32.EXE"
  /1 D:\Program Files\IBM\OnDemand32\LOCALE\ENU
```

9. Click OK.

The next time that you start the administrative client, Windows uses the start up parameters and values that you specified.

Examples

The following program properties could be used to set up the Administrator so that when a person starts the program, the name `Customer Service` is displayed in the title bar of the administrator window.

```
"D:\Program Files\IBM\OnDemand32\ARSADM32.EXE"
  /1 D:\Program Files\IBM\OnDemand32\LOCALE\ENU /T "Customer Service"
```

The following program properties could be used to set up the Administrator so that when a person starts the program, OnDemand displays only users and folders in the navigator pane. When the user selects folders, OnDemand displays only the Permissions and Field Information pages.

```
"D:\Program Files\IBM\OnDemand32\ARSADM32.EXE"
  /1 D:\Program Files\IBM\OnDemand32\LOCALE\ENU /I /D 2
```

The following program properties could be used to set up the Administrator so that when a person starts the program, OnDemand attempts to log on to the specified server, without displaying the Logon dialog box. If the logon attempt is not successful, then OnDemand opens the Logon dialog box to allow the user to log on to the system.

```
"D:\Program Files\IBM\OnDemand32\ARSADM32.EXE"
  /1 D:\Program Files\IBM\OnDemand32\LOCALE\ENU /S peak /U admin /P pw
```

Appendix G. Accessibility features

The OnDemand product includes a number of features that make the product more accessible for people with disabilities. These features include:

- Features that facilitate keyboard input and navigation
- Features that enhance display properties
- Options for audio and visual alert cues
- Compatibility with assistive technologies
- Compatibility with accessibility features of the operating system
- Accessible documentation formats

Keyboard input and navigation

Keyboard input

The OnDemand clients can be operated using only the keyboard. Menu items and controls provide access keys that allow users to activate a control or select a menu item directly from the keyboard. These keys are self-documenting, in that the access keys are underlined on the control or menu where they appear.

Keyboard focus

In Windows-based systems, the position of the keyboard focus is highlighted, indicating which area of the window is active and where the user's keystrokes will have an effect.

Features for accessible display

The OnDemand clients have a number of features that enhance the user interface and improve accessibility for users with low vision. These accessibility enhancements include support for high-contrast settings and customizable font properties.

High-contrast mode

The OnDemand clients support the high-contrast-mode option that is provided by the operating system. This feature assists users who require a higher degree of contrast between background and foreground colors.

Font settings

In Windows-based systems, the user can specify display settings that determine the color, size, and font for the text in menus and dialog windows. The OnDemand client allows the user to select the font for the document list.

Non-dependence on color

Users do not need to distinguish between colors in order to use any of the functions in this product.

© Copyright IBM Corp. 2001, 2002

Alternative alert cues

In Windows-based systems, the SoundSentry feature can be used to provide visual feedback for general application and system alerts such as warning beeps.

Compatibility with assistive technologies

The OnDemand clients are compatible with screen reader applications such as Narrator and Via Voice. The OnDemand clients have the properties required for these accessibility applications to make onscreen information available to blind users.

Accessible documentation

Documentation for the OnDemand product is available in HTML format. This allows users to view documentation according to the display preferences set in their browsers. It also allows the use of screen readers and other assistive technologies.

Appendix H. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
500 Columbus Avenue
Thornwood, NY 10594
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chrome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MECHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.
Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

© Copyright IBM Corp. 2001, 2002

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Software Interoperability Coordinator
3605 Highway 52 N
Rochester, MN 55901-7829
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

Advanced Function Presentation, AFP, AIX, DB2 Universal Database, IBM, iSeries, Operating System/400, OS/390, OS/400, and Redbooks are trademarks of International Business Machines Corporation in the United States, other countries, or both.

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries or both.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Lotus is a trademark of Lotus Development Corporation in the United States, other countries, or both.

UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

Windows and Windows NT are trademarks of Microsoft® Corporation in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

Index

Numerics

109 Cache Expiration 19
110 Cache Migration 18
124 Filesystem Statistics 18, 19
128 ApplGrp Segment Expire 5
14 DB Info Exported 7
166 ApplGrp Segment Export 7
168 ApplGrp Segment Not Available 9
171 CD-ROM Manifest 129
196 Cache Migration 19
197 Cache Migration 18
201 Concurrent Licenses 128
202 Concurrent Licenses 128
87 ApplGrp Load (System Migration) 7

A

accessibility 185
adding
  application groups 73
  applications 73
  folders 73
  groups 68
  primary storage nodes 72
  printers 71
  reports 73
  server printers 71
  servers 30
  storage nodes 72
  storage sets 72
  users 65
  users to a group 69
administrative client
  adding servers 30
  assistants and wizards 36
  changing a password 31
  hardware requirements 25
  installing 25
  logging on a server 29
  passwords 31
  program properties 181
  report wizard 36
  software requirements 25
  start up parameters 181
  starting 29
  using 23, 29
  wizards and assistants 36
administrators
  application group/folder administrator 44
  folder administrator 44
  system administrator 44
  user administrator 44
AFP data
  indexing information 82
  viewing information 81
annotation permission 50
APIs
  ARSDATE 149
  ARSDOC 153
  ARSLOAD 171
APIs *(continued)*
  how to read a syntax diagram 147
  reference 145
application group/folder administrator 44, 45
application groups
  access permission 50
  adding 73, 74
  administrator 44
  administrator permission 50
  annotation permission 50
  concepts 47
  defining 73, 74, 98, 105, 113
  document permission 50
  documents
    maintaining 18, 19
  expiration processing 5, 10, 19
  fields
    defining 77, 78
  FORMS parameter 173, 174
  importing migrated index data 10
  importing migrated indexes 9
  index data
    maintaining 5, 7, 9
  logical views permission 50
  maintenance 5, 7, 9, 18, 19
  messages 5, 7, 9, 18, 19
  migration processing 7, 9, 10, 18
  overview 47
  permissions 50, 51, 76
  query restriction 50
  storage management 76
Application Program Interfaces (APIs) 145
applications
  adding 73, 80
  AFP data 81, 82
  concepts 47
  DATASET parameter 173, 174
  defining 73, 80, 98, 106, 114
  fields
    defining 84
  indexes
    defining 84
  indexing parameters 83, 84
  load information 85
  overview 47
  print options 87
  triggers 83
archive storage manager
  migration processing 10, 13
ARSDATE
  reference 149
ARSDOC
  reference 153
  sorting data 164
ARSLOAD
  automating 179
  identifying the application to load 173
  reference 171
  running 109, 117
  starting 179
assistants and wizards 36

© Copyright IBM Corp. 2001, 2002

authority
create application groups 45
create folders 45
create groups 45
create users 45

B

backup and recovery
database 103
reports and resources 22
backup considerations 21

C

CD-ROM creation manifest
messages 129
reviewing 129
commands
ADDRPTOND 133
ENDMONOND 133
FNDKEYOND 133
PRTRPTOND 133
PRTTXTOND 133
RMVRPTOND 133
STRASMOND 133
STRIMPOND 133
STRMONOND 133
concurrent users
monitoring 128
conventions
group name 46
user name 43
copying items 35
create application groups authority 45
create folders authority 45
create groups authority 45
create users authority 45

D

data
adding 153
deleting 5, 18, 153
deleting a report 123
estimating storage space for 98
expiration processing 5, 18
importing migrated indexes 9
loading 97, 105, 113, 171
migration processing 7, 9, 17
printing 153
restarting a load process 121
retrieving 153
sorting with ARSDOC 164
unloading 123
updating 153
data loading 179
database
adding index data 101
administration 1
backup 103
concepts 3
deleting index data 5, 9
expiration processing 5, 9
fields
defining 77, 78
database *(continued)*
importing migrated index data 9
maintenance 5
migration processing 7, 9
segment table 101
updating 153
database manager
concepts 3
DATASET parameter
applications 173, 174
ARSLOAD program 173, 174
date
formats of 149
internal format 149
obtaining OnDemand internal format 149
OnDemand internal format 149
DB2
concepts 3
database 3
importing migrated index data 11
indexes 4
instances 3
journals 4
systems 3
tables 4
defining
application groups 73
applications 73
folders 73
groups 68
primary storage nodes 72
printers 71
reports 73
server printers 71
storage nodes 72
storage sets 72
users 65
deleting a report 123
deleting index data 5, 9
disability 185
document permission 50
drag and drop operations 36

E

estimating storage space 98
expiration of passwords 32
expiration processing 5, 9
expired passwords 31
exporting items 36

F

facsimile parameters 142
facsimile user exit program 178
fax parameters 142
fax user exit program 178
fields
application group
defining 77, 78
database
defining 77, 78
defining 77, 78, 89, 91
folder
defining 89, 91

fields *(continued)*
indexing
defining 84
fields permission 50
file size 95, 105, 113, 177
FILESTAT error codes 130
folder administrator 44
folders
access permission 50
adding 73, 87
administrator 44
administrator permission 50
concepts 48
defining 73, 87, 107, 115
fields
defining 89, 91
fields permission 50
named queries permission 50
overview 48
permissions 50, 51, 88
FORMS parameter
application groups 173, 174
ARSLOAD program 173, 174

G

generic indexer
parameter file 108, 116
group administration 68
group owner 46
groups
adding 68
adding users 69
assigning users 69
concepts 46
defining 68
naming 46
overview 46
owner 46
permissions 51

H

how to read a syntax diagram 147

I

image files
index data 108
loading 105
importing migrated index data 9
inactivity time out 32
index data
about 4
adding 153
creating 108, 116
deleting 5, 9, 123, 153
expiration processing 5, 9
fields
defining 84
generic indexer 108, 116
importing 9
indexes
defining 84
load information 85
loading 97, 105, 113, 171
index data *(continued)*
migration processing 7, 9
restarting a load process 121
retrieving 153
triggers 83
updating 153
indexes
defining 84
installing the administrative client 25
instances 3
internal date format 149

J

journal receivers 4

K

keyboard 185

L

licenses
monitoring 128
limitations
file size 95, 105, 113, 177
size of input files 95, 105, 113, 177
loading data 97, 105, 113, 179
loading image files 105
loading index data 97, 105, 113
loading reports 97
loading user-defined data 113
logging on a server 31
logical views permission 50
login processing 33

M

managing
servers 133
manifest
reviewing CD-ROM contents 129
maximum password age 32
messages
109 Cache Expiration 19
110 Cache Migration 18
124 Filesystem Statistics 18, 19
128 ApplGrp Segment Expire 5
14 DB Info Exported 7
166 ApplGrp Segment Export 7
168 ApplGrp Segment Not Available 9
171 CD-ROM Manifest 129
196 Cache Migration 19
197 Cache Migration 18
201 Concurrent Licenses 128
202 Concurrent Licenses 128
87 ApplGrp Load (System Migration) 7
application groups 5, 7, 9, 18, 19
CD-ROM manifests 129
errors and alerts 127
expiration processing 5, 19
licenses 128
migration processing 7, 9, 18
OnDemand errors and alerts 127
System Log 127

migration processing 7, 9
minimum password length 32
monitoring
  concurrent users 128
  licenses 128
  users 128

N

named queries permission 50
naming groups 46
naming users 43

O

obtaining OnDemand internal date format 149
OnDemand
  date formats 149
  internal date format 149
OS/400 printer file 141
output queue monitor user exit program 177

P

password 33, 175
password length 31
passwords 31, 32
permissions
  about 49
  application groups 50, 76
  concepts 49
  examples 51
  folders 50, 88
  overview 49
  strategies 53
primary storage nodes
  adding 72
  concepts 46
  defining 72
  overview 46
print parameters 141
printer file 141
printers
  adding 71
  concepts 46
  defining 71
  options 87
  overview 46
printing documents 141, 153
program properties 181

Q

query restriction 50

R

receivers, journal 4
recovery considerations 22
report wizard 36
reports
  adding 73, 153
  backup and recovery 22
  defining 73
  deleting 123, 153
reports *(continued)*
  estimating storage space for 98
  loading 97, 171
  printing 153
  restarting a load process 121
  retrieving 153
  unloading 123
  updating 153
requirements
  hardware 25
  software 25
resources 102
restarting a load process 121
reviewing a CD-ROM creation manifest 129

S

security and user administration 65, 68
segment table 101
server fax parameters 142
server fax user exit program 178
server print parameters 141
server printers
  adding 71
  concepts 46
  defining 71
  options 87
  overview 46
servers
  adding 30
  adding items 35
  copying items 35
  drag and drop operations 36
  exporting items 36
  logging on 31
  managing 133
  system parameters 34
  updating 30
size of input files 95, 105, 113, 177
sorting data with ARSDOC 164
start up parameters 181
starting the administrative client 29
storage
  administration 13, 72
  for reports 98
  management 13, 76
storage nodes
  adding 72
  concepts 46
  defining 72
  loading data 101
  overview 46
storage objects
  loading 101
storage sets
  adding 72
  application group
    assigning to 76
  assigning to an application group 76
  concepts 46
  defining 72
  loading data 101
  migration processing 10
  overview 46
  System Migration application group 10
STRIMPOND command 11

syntax diagram
how to read 147
system administrator 44
System Log messages
109 Cache Expiration 19
110 Cache Migration 18
124 Filesystem Statistics 18, 19
128 ApplGrp Segment Expire 5
14 DB Info Exported 7
166 ApplGrp Segment Export 7
168 ApplGrp Segment Not Available 9
171 CD-ROM Manifest 129
196 Cache Migration 19
197 Cache Migration 18
201 Concurrent Licenses 128
202 Concurrent Licenses 128
87 ApplGrp Load (System Migration) 7
application groups 5, 7, 9, 18, 19
CD-ROM manifests 129
enabling 33
expiration processing 5, 19
loading data 102, 111, 119
migration processing 7, 9, 18
monitoring 127
user exit programs 33
user messages 128
System Log user exit program
migration processing 9
System Migration application group 10
system parameters
about 32
inactivity time out 32
login processing 33
maximum password age 32
minimum password length 32
password 33
password age 32
password expiration 32
password length 32
passwords 32
setting 34
System Log 33
time out 32
user exit logging 33
userid 33
systems 3

T

tables 4
TIF files
loading 105
time out 32
triggers
defining 83

U

unloading a report 123
update servers dialog box 30
user 44, 45
user administration 65
user administrator 44, 45
user exits
facsimile user exit program 178
fax user exit program 178
user exits *(continued)*
migration processing 9
output queue monitor user exit program 177
reference 177
saving messages in the System Log 33
server fax user exit program 178
System Log 9
user-defined data
index data 116
loading 113
userid 33, 175
users
adding 65
adding to a group 69
application group/folder administrator 44
assigning to a group 69
authority 45
changing a password 31
concepts 43
defining 65
folder administrator 44
licenses 128
monitoring 128
naming 43
overview 43
passwords 31
permissions 49, 51
system administrator 44
types 44
user 44
user administrator 44
using the administrative client 23, 29

W

wizards and assistants 36



Program Number: 5722–RD1



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

SC27-1161-01



Spine information:



IBM Content Manager OnDemand
for iSeries Common Server

Version 5 Release 2